

An analysis of contractual and transactional aspects of a teleradiology process

Citation for published version (APA):

Vonk, J., Wang, T., Grefen, P. W. P. J., & Swennenhuis, M. (2008). *An analysis of contractual and transactional aspects of a teleradiology process*. (BETA publicatie : working papers; Vol. 263). Technische Universiteit Eindhoven.

Document status and date:

Published: 01/01/2008

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

An analysis of contractual and transactional aspects of a teleradiology process

Technische Universiteit Eindhoven,
department of Technology Management,
Information Systems group:
ir. Jochem Vonk,
M.Sc Ting Wang,
Prof. Dr. ir. Paul Grefen

Topicus HealthCare:
Marcel Swennenhuis

Table of Contents

1	Introduction.....	5
1.1	The Healthcare Domain	5
1.2	The XTC Project	5
1.3	Case Study: Teleradiology	7
1.4	Structure of this report	8
2	The Teleradiology Process.....	9
2.1	The Main Teleradiology Process	9
2.2	Scan Interpretation Subprocesses.....	11
3	Service Oriented View	14
4	Agreements	19
4.1	Service Requirements	19
4.2	Service Offers	20
4.3	Established e-Contracts.....	23
5	Introducing XTC in the Teleradiology process.....	29
5.1	Transactional Quality of Service.....	29
5.1.1	Fluency.....	30
5.1.2	Interference	33
5.1.3	Alternation	34
5.1.4	Transparency	37
5.2	Abstract Transactional Constructs	38
5.3	Mapping	42
6	Summary and Conclusions.....	46
7	References.....	48
Appendix A	Summary of the processes.....	51
Appendix A.1	The Scan Interpretation Process.....	51
Appendix A.2	The Main Teleradiology Process	52

List of Figures

Figure 1: ATC and TxQoS context.....	6
Figure 2: Main Teleradiology Process	10
Figure 3: Scan Interpretation process.....	12
Figure 4: Service Oriented View on the teleradiology process.....	14
Figure 5: Scan Interpreter invokes three services.	16
Figure 6: Scan Interpreter invokes one service	16
Figure 7: Scan interpreter invokes two services	17
Figure 8: External view of the scan interpretation process	21
Figure 9: Mapping internal level to external level process	22
Figure 10: From requirements and offers to e-Contract.....	24
Figure 11: Transactional External View of the Scan Interpretation process.....	32
Figure 12: Interference External View of the Scan Interpretation process	33
Figure 13: Alternate External View of the Scan Interpretation process	35
Figure 14: Compensations for the Main Teleradiology process	40
Figure 15: Teleradiology ATC Tree	42
Figure 16: Internal Scan Interpretation process with Compensation and Alternative	43
Figure 17: Internal compensation versus external compensation	44
Figure 18: External view of the scan interpretation process	52
Figure 19: Global Process view	53

1 Introduction

This report presents the results of an extensive case study conducted in the teleradiology area. Focus of the case study is the analysis of contractual and transactional aspects of a teleradiology process. Thereby applying and validating the concepts and ideas developed in the XTC project within this process. The following subsections present the context in which this case study has been performed.

1.1 *The Healthcare Domain*

Processes in the healthcare domain are often very complex. They contain numerous subprocesses and activities covering multiple departments in a hospital and/or hospitals and/or other organizations. Included subprocesses range from relatively simple daily routine tasks to complex medical procedures, see for example [VW+08]. As healthcare processes deal with patients, reliability of such processes is even more important than it is in many other complex business processes.

Recent years show a trend in increasing numbers of collaborations between organizations, ranging from simple outsourcing of subprocesses to complex business network processes [GMK07]. With the developments in information technology, the same trend is starting to appear in the healthcare domain as well. Already existing examples are teleradiology [Sch06, Ran04] (the focus of this report) and telemedicine (the modern version of 'in absentia healthcare'). Through this trend, the main focus shifts from the departmental functional view (as it is now) towards a cooperation view involving healthcare providers from different medical disciplines that can exist within the organization/hospital (the departments) or that exist in other hospitals and/or organizations (as it will be in the future).

From a patient point of view, the healthcare processes should be transparent (up to a certain level) and the quality of the process should be known (and measurable). More and more patients want to be able to choose their healthcare provider themselves, basing their choice on the quality of the healthcare process they will get involved in. This means that healthcare providers will start to focus more on improving the quality attributes (like performance indicators [PIZ08]) of their healthcare processes, so that they are able to distinguish themselves from other healthcare providers. A base-set of quality attributes (quality indicators) is given in [PIZ08]. An example of a more advanced quality attribute for a healthcare process can be the reliability of such a process, e.g., what is the chance that a planned medical procedure will be cancelled and what will happen if a cancellation does actually happen.

1.2 *The XTC Project*

Increasing process execution reliability is the focus of the XTC project, carried out at Eindhoven University of Technology and Tilburg University. XTC is an acronym for eXecution of Transactional Contracted electronic services. Within the XTC project, the business transaction framework (BTF) has been developed, which consists of concepts, techniques, and methods that, when applied to a

business process, increase its execution reliability, increases the flexibility in choosing collaboration parties and facilitates the explicit specification of the collaboration in agreements (e-Contracts), with the focus on business process execution reliability.

Two key elements of the BTF are abstract transactional constructs (ATC) and transactional quality of service (TxQoS). ATCs provide flexible reliability in the execution of processes by applying transaction management techniques [WGV06]. Cooperation between organizations usually takes place using services. A Service is a business function (implemented and/or supported by a business process) that is offered by an organization to other organizations. The agreements related to a cooperation are stated in an electronic contract. Agreements specifically related to services are contained in a service level agreement (SLA), which is part of the contract. The SLA in turn can contain a specific section covering transactional issues related to the service, called the transactional service level agreement (TxSLA). The TxSLA contains transactional quality of service (TxQoS) specifications that together determine the overall reliability agreements (related to that service) between the involved organizations.

Figure 1 shows the relation between the different concepts¹. Processes are supported by ATCs. TxQoS is specified for Services and contained in a TxSLA. Services contain processes (while processes can contain or make use of services as well). As the actual enactment of a service is performed through the process execution it contains, the TxQoS of the service is guaranteed by the ATCs corresponding to the process, i.e., TxQoS is mapped to ATCs and vice versa. A service is governed by a SLA. The SLA itself is contained within a contract, and the TxSLA is part of the SLA.

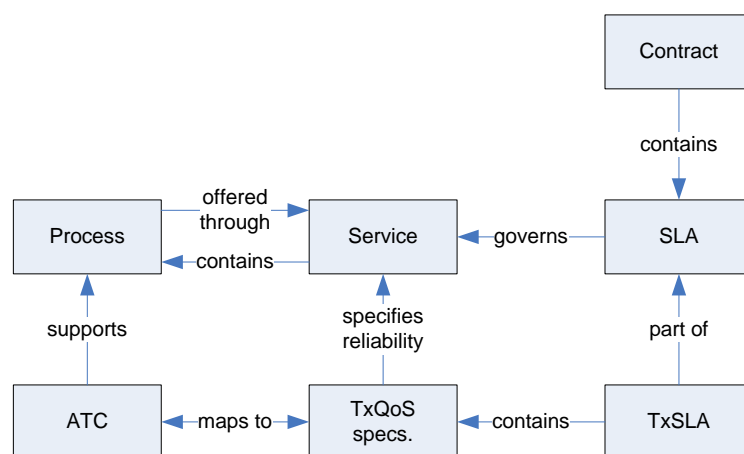


Figure 1: ATC and TxQoS context

¹ To keep Figure 1 simple, the cardinalities on the relations between the different concepts are omitted. A more detailed version of this diagram can be found in WGV08 and VWG07.

To evaluate the business transaction framework, we have conducted two case studies in the healthcare domain. The first case study considers the cardiothoracic surgery (CTS) process of the Catharina Hospital in Eindhoven, and is described in [VW+08]. This case study focuses more on a detailed description (modeling) of the process. The XTC concepts and ideas are applied to the process and the existing contracts (or SLAs) to illustrate how the process execution reliability can be improved.

This report describes the second case study conducted in the teleradiology area. The focus of this second case study is on process reengineering; to illustrate how the process specification changes and what the agreements between the involved parties should contain when the XTC ideas and concepts are applied to improve the execution reliability of this process.

1.3 Case Study: Teleradiology

The starting point for this case study is an existing teleradiology process [Swe07]. Teleradiology concerns the acquisition and interpretation of radiology scans, for example CT scans or X-ray scans. However, the interpretation of the scans takes place in a different location than the scan acquisition itself. The process models representing the existing Teleradiology process are presented in Section 2. By taking a service-oriented view on the process, it is possible to explicitly identify the services involved in the process. Through the identification of the services and the parties that provide these services, it becomes possible to determine which agreements are required and between which parties such an agreement should be made. In the agreements quantitative and qualitative characteristics (or non-functional characteristics) of the services can be specified. Understanding the qualitative service characteristics (for example the reliability of the service execution) and the possibilities in changing them (for example decreasing the chance of exceptions in the service execution) can provide a competitive advantage to a service provider, by offering better qualitative characteristics for the service on offer compared to similar services offered by other health care providers (teleradiology organizations in this case).

Through the application of the developed business transaction framework the teleradiology process is complemented with structured exception handling, concurrency control, and visibility control. This increases the overall quality and reliability of the process in the presence of occurring exceptions, that are unavoidable in the health care domain as it is highly dynamic and unpredictable. A service oriented view, together with concepts of e-contracting increases the flexibility in the end-to-end teleradiology process by increasing the choices for collaboration parties and collaboration constellations. This readies the teleradiology domain for the future in which services and especially the quality of services aspects becomes more and more important.

1.4 Structure of this report







The teleradiology process as it currently exists is described and illustrated in the next section. Section 3 introduces the service oriented view of the teleradiology process and presents a number of possible collaboration constellations. The agreements that have to be made between collaborating parties and the process of reaching agreements, by matching service offers and requirements and establishing an electronic contract, is covered in Section 4. Improvements to the teleradiology process with the concepts and ideas developed in the XTC project is presented in Section 5. First the FIAT framework is introduced covering the improvements specified in the agreements. Second the ATC concept is explained and illustrated by applying it to the teleradiology thereby supporting its execution through transaction management and hence improving the reliability of the process execution. Third, the mapping between the FIAT framework and the supporting ATCs is described. The report is concluded in Section 6, with a summary and conclusions.

2 The Teleradiology Process

The teleradiology process, as described in this report, is applicable in a number of settings: it can be used in the acquiring of scans within a hospital, but also for acquiring scans in a private clinic or mobile scan acquisition facility. The interpretation of the acquired scans, however, is performed by a different organization (or department) at a different location (hence the tele in teleradiology). The main teleradiology process starts from the moment an order for a scan is received and ends after payment has been made and an analysis on the performance of the scan acquisition and interpretation has been made.

The modeling notation used is that of electronic process chains (EPC), the legend of which is shown in Table 1 below [Sch00, KNS91].

Table 1: Legend for EPC diagrams

	Event: a status change. Can be caused by performing a function, but can also be caused by an external factor (e.g., time).
	Function: c.f. process activity or task
	Organizational Unit: which organizational unit or role performs the associated function.
	eXclusive Or-Split or eXclusive Or-Join. In case of a split: when the incoming arrow becomes 'true' (the previous event happened or the previous function was performed), only one of the outgoing arrows is 'followed'. In case of a join: when one of the incoming arrows becomes 'true', the outgoing arrow is 'followed'.
	Or-Split or Or-Join. Similar to the exclusive or, except that more than one outgoing arrows (split) can be followed, or more than one of the incoming arrows (join) can be 'true'.
	And-Split or And-Join. In case of a split, all outgoing arrows are followed in case the incoming arrow becomes 'true'. In case of a join: all incoming arrows must be 'true' before the outgoing arrow is followed.

2.1 The Main Teleradiology Process

The main teleradiology process is shown in Figure 2. The process starts as soon as a scan order is received. This order can come from a hospital (i.e., a specialist in the hospital) or from a person

directly (a new trend in radiology, in which persons contact a radiology clinic directly to make scans without a direct medical need or referral, see for example Pre08 and Roy08). The resulting scan report is returned to the person that requested the scan.

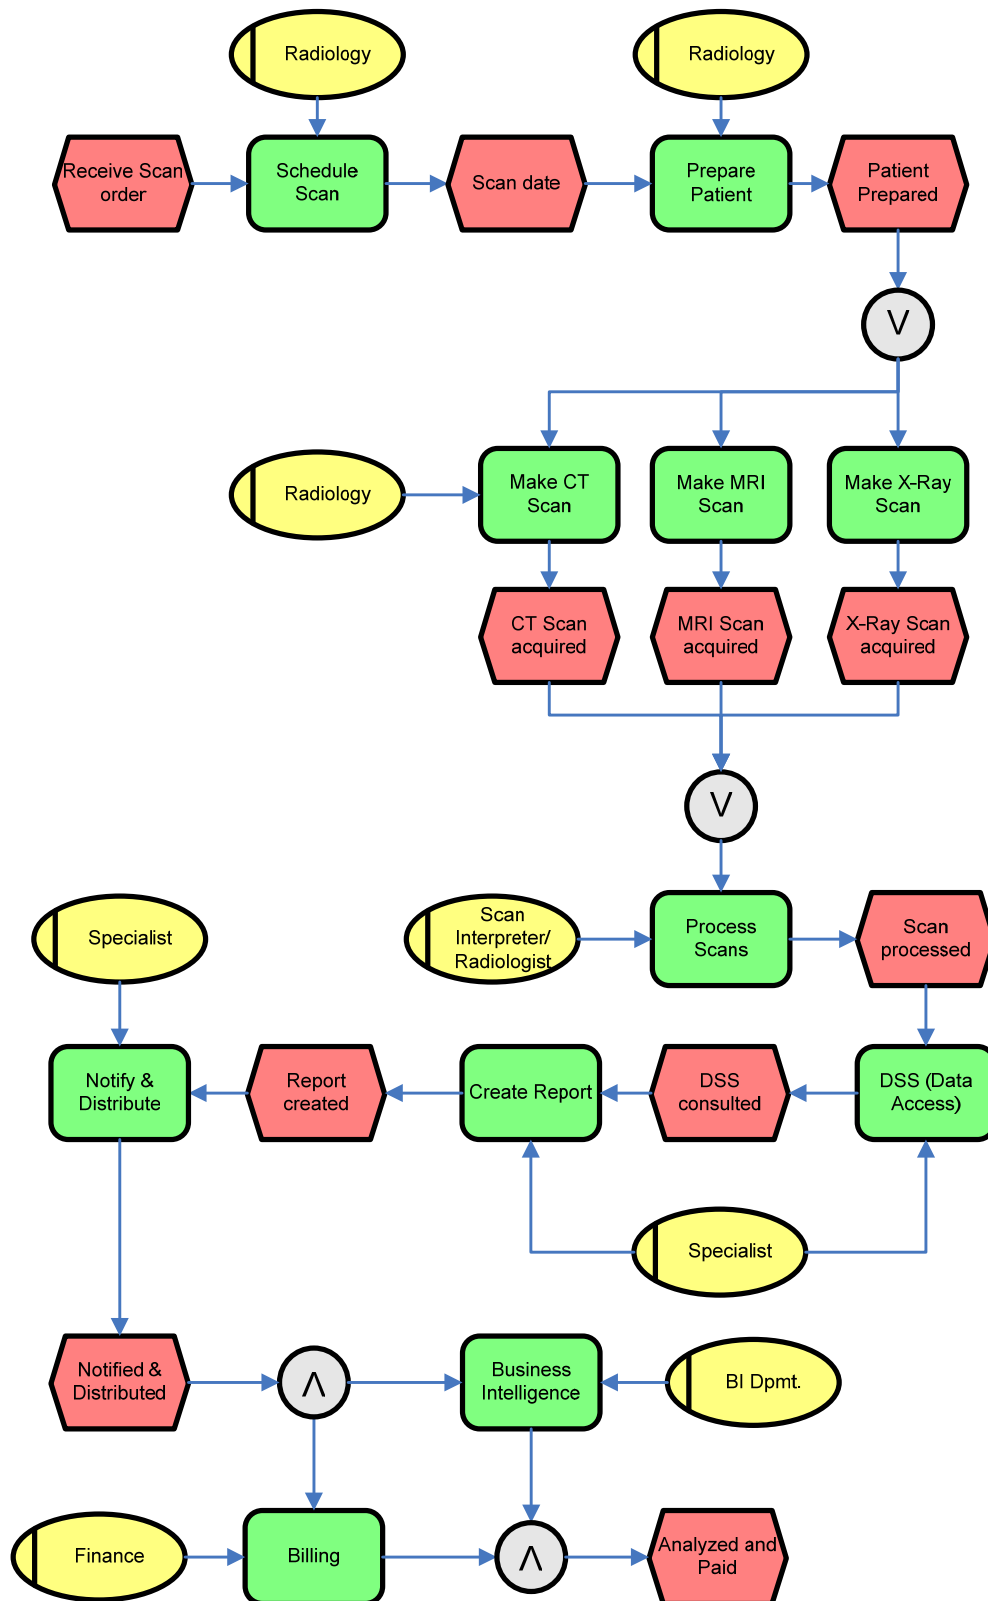


Figure 2: Main Teleradiology Process

After receiving the order, the scan acquisition needs to be scheduled (and the patient informed of the date). At the scheduled date, the patient is prepared so that he can undergo the scan acquisition activities. Examples of the different types of scans are: a computed tomography (CT) scan, a magnetic resonance imaging (MRI) scan, or an X-Ray scan. Which scans are required is specified in the scan order. When the required scans are made, they need to be processed (analyzed and/or interpreted). This part of the process (the '*Process Scans*' subprocess in Figure 2) is outsourced to either a dedicated organization specialized in performing scan interpretations, or to individual specialists that have the facilities to perform scan interpretations themselves. Details of the outsourced process are given in the next subsection.²

The result of the outsourced process (the processing of scans) is a report that contains the result of the scan analysis. These results are stored in an information system so that they can be used by (advanced) clinical applications. In this process, the result of the scan interpretation is used in a decision support system (DSS) to help in deciding what the best follow-up activities for a patient are. A report is created that contains this decision together with the argumentation, including (a summary of) the scan interpretation results. The person that ordered the scan (i.e., the person that initiated this process) is notified that the report is ready and a copy of this report is sent to him. At this stage, the process part related to the scan acquisition and interpretation is finished and only the payment handling and an analysis of the process (executed so far) remains, after which the process is completed.

2.2 Scan Interpretation Subprocesses

The outsourced scan interpretation process is shown in Figure 3. When a scan is received, it is scheduled to be interpreted. At the scheduled time, the scan is retrieved and analyzed. The analysis is performed by a specialist (radiologist) who looks at the scan to discover anomalies and at the same time dictates how he interprets the scan. Recently, software recognition is being introduced to be used in the direct transformation of the voice dictation into a digital textual format, see for example [Nua08]. Alternatively, the voice dictation has to be transcribed into a digital textual format. In case the specialist dictates in a different language than the language in which the scan analysis report is to be delivered, a translation of the digital textual scan interpretation is necessary. The scan interpretation, together with the optional transcription and translation, forms the report created by the specialist describing his findings about the scan.³ If the report is written in a language not native to the

² In the teleradiology process a number of different scans can be required for one scan order. As modeled in Figure 2, the different scans are collected and processed together in a batch. Other possibilities exist, in which, for example, all scans are processed individually after which the result must be combined so that the Decision Support System can be accessed and a diagnosis can be made.

³ Because the teleradiology process and the scan interpretation process belong to different parties, activities with the same name can occur, even though the content of these activities is most likely very different. In this example, this is the case for the 'Create Report' activities.

specialist, a native speaker verifies the language used in the report. If mistakes are found, they have to be corrected either by the specialist (radiologist) or, in case of simple syntactical mistakes, by the native speaker himself.

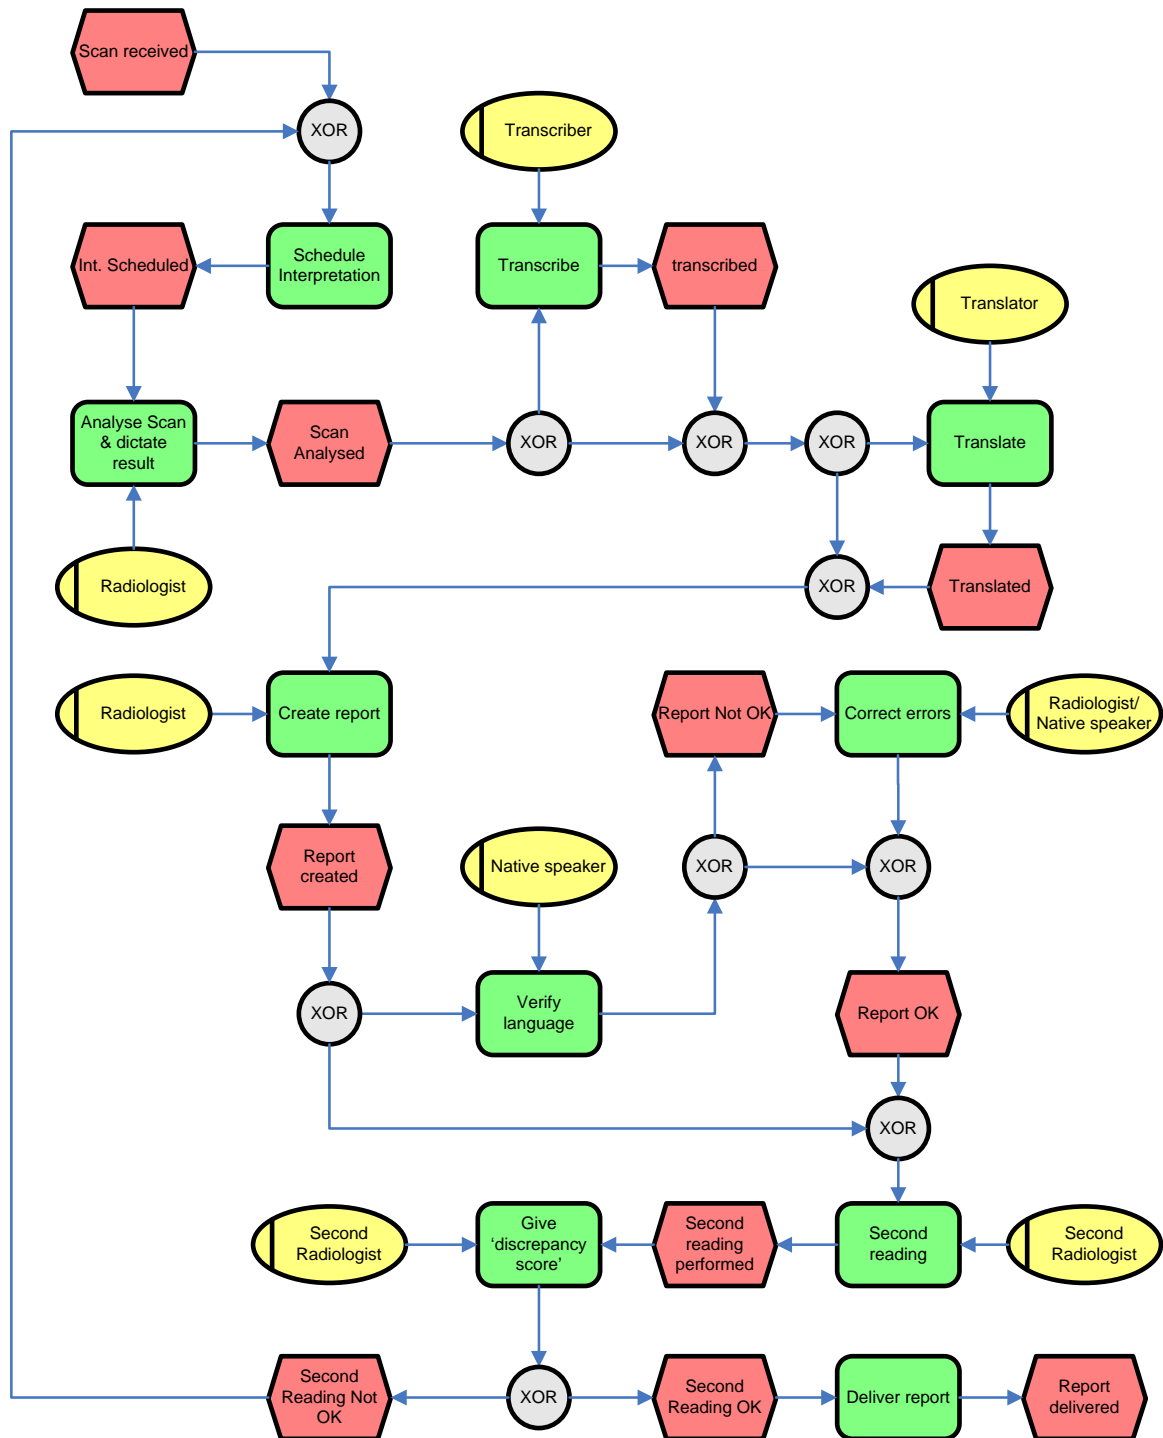


Figure 3: Scan Interpretation process

When the report is verified, a second reading takes place to further reduce the chance of analysis mistakes. In the second reading (also called double reading) the report and the scans are checked by another specialist/radiologist than the person who made the report. The person doing the second reading gives a, so called, discrepancy score that indicates whether or not the report adheres to the quality standards agreed for the scan interpretation. If it is found not to be ok, the scan analysis, including the report creation, has to be redone. If the report is determined to be ok after the second reading, the report is delivered to the requestor of the scan interpretation process and the process finishes. Billing is not part of this process because it is performed batch-wise and executed in another separate process.

The teleradiology process, including the scan interpretation subprocess, as described in this section, represents the situation as it currently exists in the teleradiology domain. Agreements are made between hospitals and clinics, hospitals and ‘scan interpreters’, and/or between clinics and ‘scan interpreters’. These agreements contain general provisions that apply to all instances of the teleradiology process. Examples of general provisions are turn around times, duration of the contract, and the maximum number of invocations of the process. However, through advancements in the information systems research field, new opportunities arise that allow for more complex collaboration constellations which can be established quicker and which are governed by e-contracts and service level agreements containing detailed specifications on qualitative service aspects. The following sections present enhancements to the teleradiology process by applying the ideas and concepts that have been developed within the XTC project. This way, it illustrates how the current teleradiology process can be changed into a future version that can be used to gain a competitive edge. More specifically, it illustrates how the process specification should change, which agreements should exist and what those agreements should contain (with a focus on process execution reliability). The Service Oriented View of the teleradiology process is presented in the next section. It facilitates the identification of involved services and parties and thereby helps in determining the possible collaboration constellation between the involved parties.

3 Service Oriented View

As briefly explained in the introduction, the XTC project aims to improve reliability in process executions (using ATCs) and to facilitate creating reliability agreements between (units of) organizations (TxSLA). The relevant concepts and their relations are shown in Figure 1.

In the current situation, as described in the previous section, only two parties are involved in executing the Teleradiology process (not counting the party that requests the scan order): one party takes care of the scan acquisition and the other party takes care of the scan interpretation. However, from the process models shown in Figure 2 and Figure 3, we can identify a number of business functions that could be performed by other organizations. These services are:

- a transcribe service,
- a translator service,
- a language verification service,
- a ‘second’ reading’ service,
- a billing service, and
- a business intelligence service.

Note that the reports are delivered in digital form, otherwise a postal service would also be included (as part of the ‘*Notify & Distribute*’ subprocess).

From the identified services and the process specification models, a service oriented view can be constructed, which illustrates the possible interactions between the parties (and their services) involved. This service oriented view is shown in Figure 4 below.

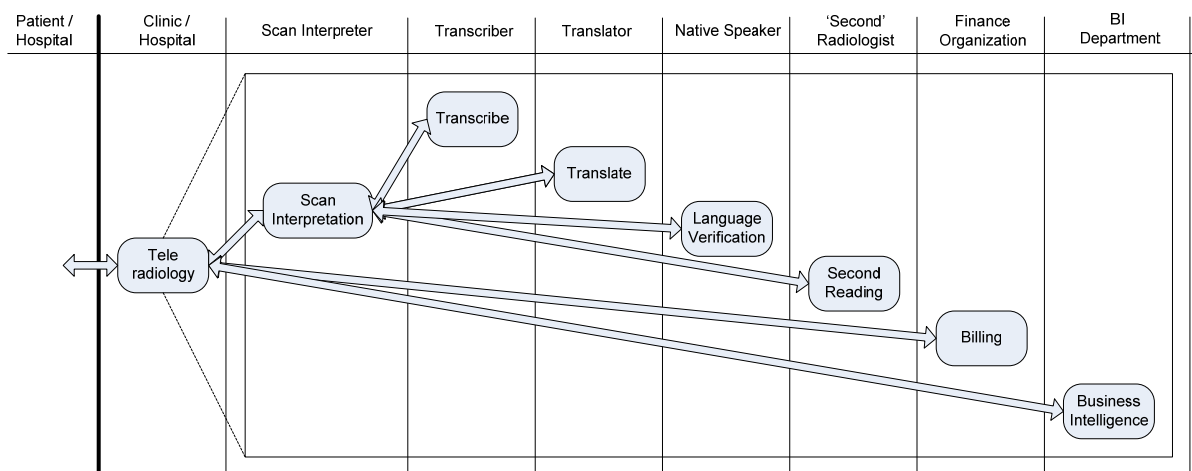


Figure 4: Service Oriented View on the teleradiology process

The columns in the figure represent the parties involved in the teleradiology process and the rounded rectangles represent services that are provided by the involved parties. The arrows indicate the possible interactions.

From the point of view of the person initiating the process (the person giving the scan order), only one service exists: the teleradiology service⁴. This service is offered by specialized clinics (to patients and/or hospitals) or hospitals (to departments of the same hospital or to other hospitals).

The clinic or radiology department in the hospital, in turn, makes use of the scan interpretation service offered by other organization(s) to process the acquired scans. In this case, the clinic can play two roles in this process; one role is the provider role in which the clinic provides the teleradiology service, the other role is the consumer role in which the clinic consumes the scan interpretation service offered by another organization.

The organization providing the scan interpretation service (the scan interpreter) makes use of a number of other services to be able to produce a transcribed, translated, verified, and double checked scan analysis report⁵. However, instead of calling all these services directly, the same report can be produced by invoking these services in a different interaction configuration. The interaction configuration (or collaboration constellation) determines the way the involved organizations can collaborate, i.e., it determines the possible agreements that can be made, the visibility of certain service providers to other organizations in the collaboration, etc. In this scenario, the scan interpreter has the role of service consumer, requiring a transcribed, translated, and verified scan interpretation report. The following list presents the possible service interaction configurations for the scan interpreter (or different variations of the service oriented view with respect to the scan interpretation service, ignoring the second reading service as that is always called by the scan interpreter directly):

1. The scan interpreter invokes the other three services directly, as shown in Figure 5. After the scan has been interpreted (dictated) the transcribe service (1), the translate service (2), and the language verification service (3) are invoked, in that order.

⁴ Although the service is called teleradiology, from the perspective of the person initiating the scan order it is a radiology service, i.e., the fact that it is ‘tele’ is not relevant to that person.

⁵ The services invoked by the scan interpreter could also be called directly by the clinic/hospital, meaning that there is no indirection at all: the clinic/hospital invokes all services and manages the synchronization between the invoked services. However, this scenario is not very realistic and is therefore ignored in this report.

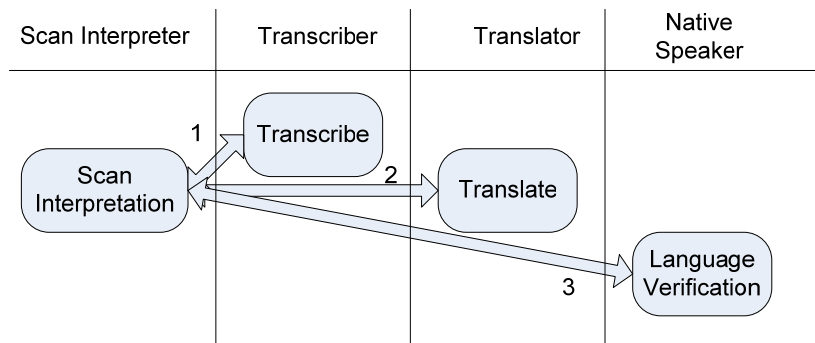


Figure 5: Scan Interpreter invokes three services.

2. The scan interpreter invokes only one service (the transcribe service), which in turn invokes the other services. Figure 6 shows the two different collaboration constellations that are possible in this case. In the top part of the figure, the transcribe service invokes the translation service (2), which, in turn, invokes the language verification service (3). In this case, the transcribe service is not aware of the language verification service being offered and executed by another organization than the translator. In the bottom part of the figure, the transcribe service invokes the translate service (2), and after receiving the result back, it invokes the language verification service (3).

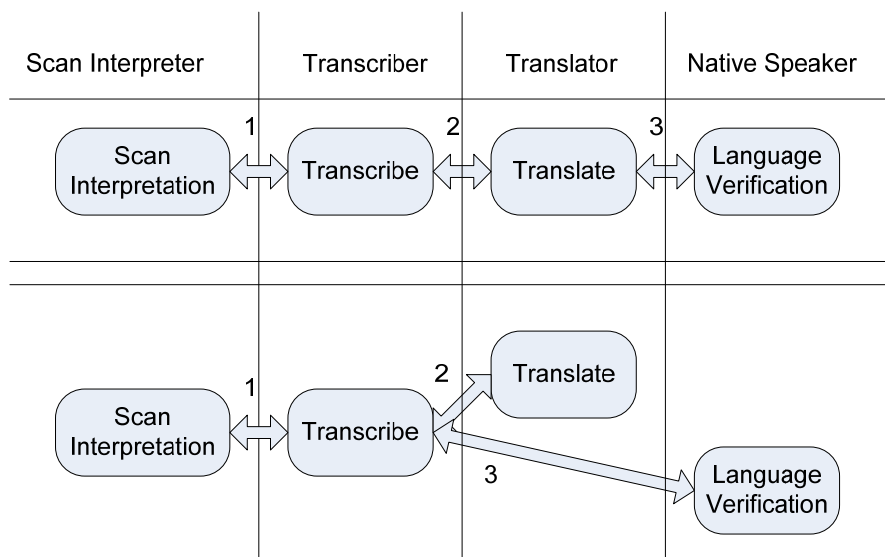


Figure 6: Scan Interpreter invokes one service

3. The scan interpreter invokes two of the three services, see Figure 7. Again, two different collaboration settings are possible. In the first setting (top part of the figure), the scan interpreter invokes the transcribe service (1) to acquire a transcribed and translated report from the dictation. For the translation, the transcribe services makes use of a translation

service (2), which is not visible to the scan interpretation service. After receiving the result back from the transcribe service, the scan interpretation service invokes the language verification service (3). The bottom part of the figure shows the setting in which the scan interpreter invokes the transcribe service (1) and receives, as a result, a transcribed report. Then, the translate service is invoked (2), passing the transcription as an input. The translate service invokes the language verification service (3), which is not seen by the scan interpreter service. The result is a verified translation, which is then returned to the invoker, i.e., the scan interpreter.

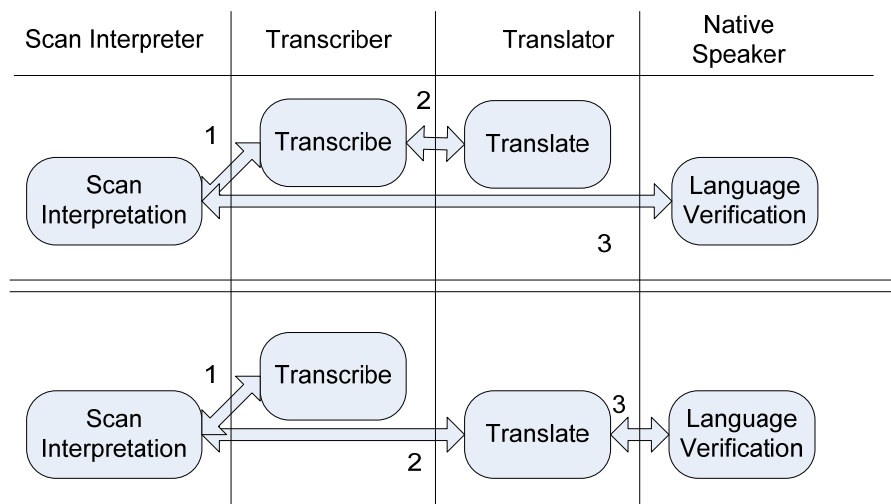


Figure 7: Scan interpreter invokes two services

Each of the collaboration constellations has its advantages and disadvantages. As it is the service consumer who can decide which organization it desires as direct collaborating parties, the advantages and disadvantages are with respect to the service consumer. The more parties (service providers) an organization directly collaborates with, the more visibility and control that organization can have, but which also leads to more complexity in contract management (see Section 4). For example, in the first setting mentioned above, the scan interpreter knows exactly who the other organizations are that he/she collaborates with. In the situation in which the scan interpreter only outsources to the transcribe service (setting 2 above), the other services (translation and verification) are not visible to the scan interpreter, and therefore the control over these services is non-existent (or limited at most) with respect to the scan interpreter.

Note that not all three services (transcribe, translate, and language verification) are required in all process instances, i.e., they can be optional. In such a situation, not all collaboration settings are even possible.

As the service oriented view makes explicit which organization interacts with which other organization(s), it becomes clear which collaborations exist in this setting and between which organizations collaborations takes place. Each collaboration (or service invocation) is setup to exchange some values, and each value exchange is accompanied by rights and obligations for each of the involved organization. To exactly know what the rights and obligations of the collaborating parties are, contracts are established between them. Specific agreements between parties that are related to the process executions are specified in service level agreements (SLAs), which is part of the contract. As each collaboration is governed by a contract, this implies that a higher number of direct collaboration partners leads to a higher number of contracts that need to be established. So more direct collaboration provides more visibility and control, but it also requires more complex contract management and more monitoring of the enactment of those contracts. The next section describes agreements, service requirements, service offers, and electronic contract in detail.

4 Agreements

Any collaboration between different organizations (or even between departments within an organization) in which the rights and obligations of the involved parties as well as the penalties in case the obligations are violated need to be explicitly stated, needs to be specified in a contract. Electronic contracting, also called e-contracting, is the approach in which contracts are specified in an electronic format that can be interpreted by an information system so that an automated support for enacting contracts becomes possible. Automatic contract establishment, including negotiations, is also considered part of e-contracting. This has become feasible due to the digitalization of requirements and offers, which is used in matching organizations and in negotiations between matched organizations to reach a consensus in the rights and obligations adhering to the e-contract. Additionally, e-contracting also introduces new business opportunities [AG04, GA02]. For example, contracts can now be established within seconds (instead of the usually long duration it takes in traditional (paper) contract establishment). Through the automated interpretation of contracts an enactment infrastructure can be setup in a fully automated way. Also, monitoring the enactment progress can be done automatically. However, the amount of monitoring and what can be monitored depends on the agreements made in the contract. More details on the electronic contracting area can be found in the [AG03], in which the 4W framework is introduced. The 4W framework considers the Who, Where, What, and hoW concepts that are part of a contract. In this report, only the What and hoW are considered, while the Who and What concepts are considered out of scope.

In case the subject of a contract concerns service invocation, the contract is usually accompanied by (or contains) a service level agreement (SLA) in which the quality aspects related to the service enactment are explicitly specified in a clear and unambiguous way. The bare minimum that is contained in a SLA is the process specification (possibly only the activities without the control flow). This process specification is used for monitoring the progress of the service execution. Note that the process specification could be a black box, in which case the service is represented to the consumer by only one activity [GLA03].

Before an e-contract can be established, the service requirements and service offers should be specified so that they can be matched, and a collaboration can be formed (possibly after negotiations).

4.1 Service Requirements

An organization that wants a service to be performed by another organization needs to specify its requirements for that service to be able to find a suitable service. Potential suitable services should be able to satisfy these requirements (if requirements are negotiable then they are less important in the consideration whether a service is potentially suitable or not). To facilitate matchmaking between the service requirements and service offers (see next subsection), the service requirements are specified in

as a service level agreement specification (although they have not been agreed upon yet as they are merely service requirements; it is a unilateral SLA-specification).

Numerous types of requirements exist; the following lists a number of the most common requirements types (the concrete value of each requirement type is service specific):

- Service description / Service goal
- (part of the) process specification
- Duration (average, maximum)
- Costs (per service invocation, per time period, per batch, etc.)
- Invocations (maximum number, time spread)
- Input and output formats
- Monitoring service progress (what, when, and how to monitor)
- Level of control (cancellation possibilities, pause possibilities, etc.)
- Reliability or Expected Failure Rate (how many expected failures per time unit)
- Recovery semantics (how failures are recovered and how long does recovery take). Note that if recovery is done by means of another process (e.g. compensation process), then the same type of requirements can be imposed on such processes as well.

The service requirements are used to search for potential collaboration parties, by matching these requirements with service offers offered by other organizations. If no matching offers are found, the service requirements can be stored in an ‘open’ repository that is accessible to other organizations. Then other organizations are able to see that some organization is searching for that specific service, so that they can create a service offer to match (or nearly match if negotiations are desired) the required service and a collaboration can be established.

4.2 Service Offers

Similar to the requirements listed above, an organization that wants to offer a service needs to specify those quality of service aspects that a potential service consumer could find important, or that could distinguish its service from similar services offered by other service providers. In order for consumer organizations to discover existing services, they are listed in some form of ‘yellow pages’, for example a UDDI repository [Udd08]. These yellow pages facilitate searching and browsing of services and are used by matchmaking systems to match service requirements with service offers. So a possible service offer can contain the same quality of service aspects for the service on offer as listed in the previous subsection for the requirements, however, the values associated with the quality aspects can differ substantially between offers and requirements. Also, as mentioned in the previous subsection, a service offer can be specifically created to match a required service. In this case, it is not necessary to store the service offer in a repository.

Organizations can distinguish their service offers from other organizations by listing more service aspects or by offering ‘better’ values for those aspects. For example, a service offer for a service with an expected failure rate of zero is better than a similar service offer but with an expected failure rate of ten. However, the first offer is only ‘better’ if expected failure rate is an important issue for the organization looking for that service.

A mandatory aspect of the e-contract offer is of course a description of the service, i.e., what does the service do/produce/result in or what is the goal of the service. This is the aspect that most likely will be used by other organizations to search for services they need. For monitoring purposes, the service is usually not just a black box service. A black box service contains merely a description of the service, without any process structure information. Services are usually ‘opened up’, to reveal more information on the process structure, required for progress monitoring, but can also be required by an organization, for example for quality control. Because an organization offering a service is usually not willing to expose its detailed process (to keep its competitive edge with respect to competitors), an abstraction of the detailed process is created, called external level view, which is then exposed in the service offer, see the 3-level framework in [GLA03]. Most of the (other) QoS specifications in the SLA are specified in relation to the exposed service abstraction.

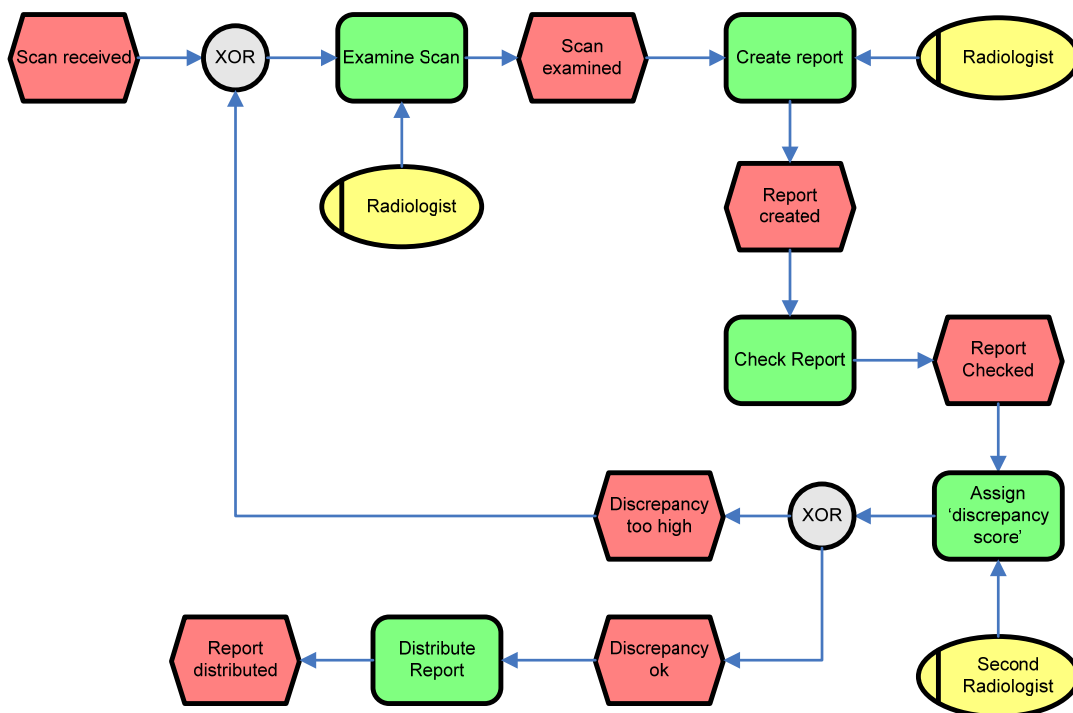


Figure 8: External view of the scan interpretation process

Figure 8 shows an example of an external view of the scan interpretation process, as specified in Section 2.2 and shown in Figure 3. This external view process consists of a number of activities and

events that are abstractions of the internal level activities. Also the role attached to activities at the external level can be an abstraction of the roles specified for the internal level process. The abstraction relation between the internal level process, as shown in Figure 3, and the external view process, shown in Figure 8, is illustrated in the figure below.

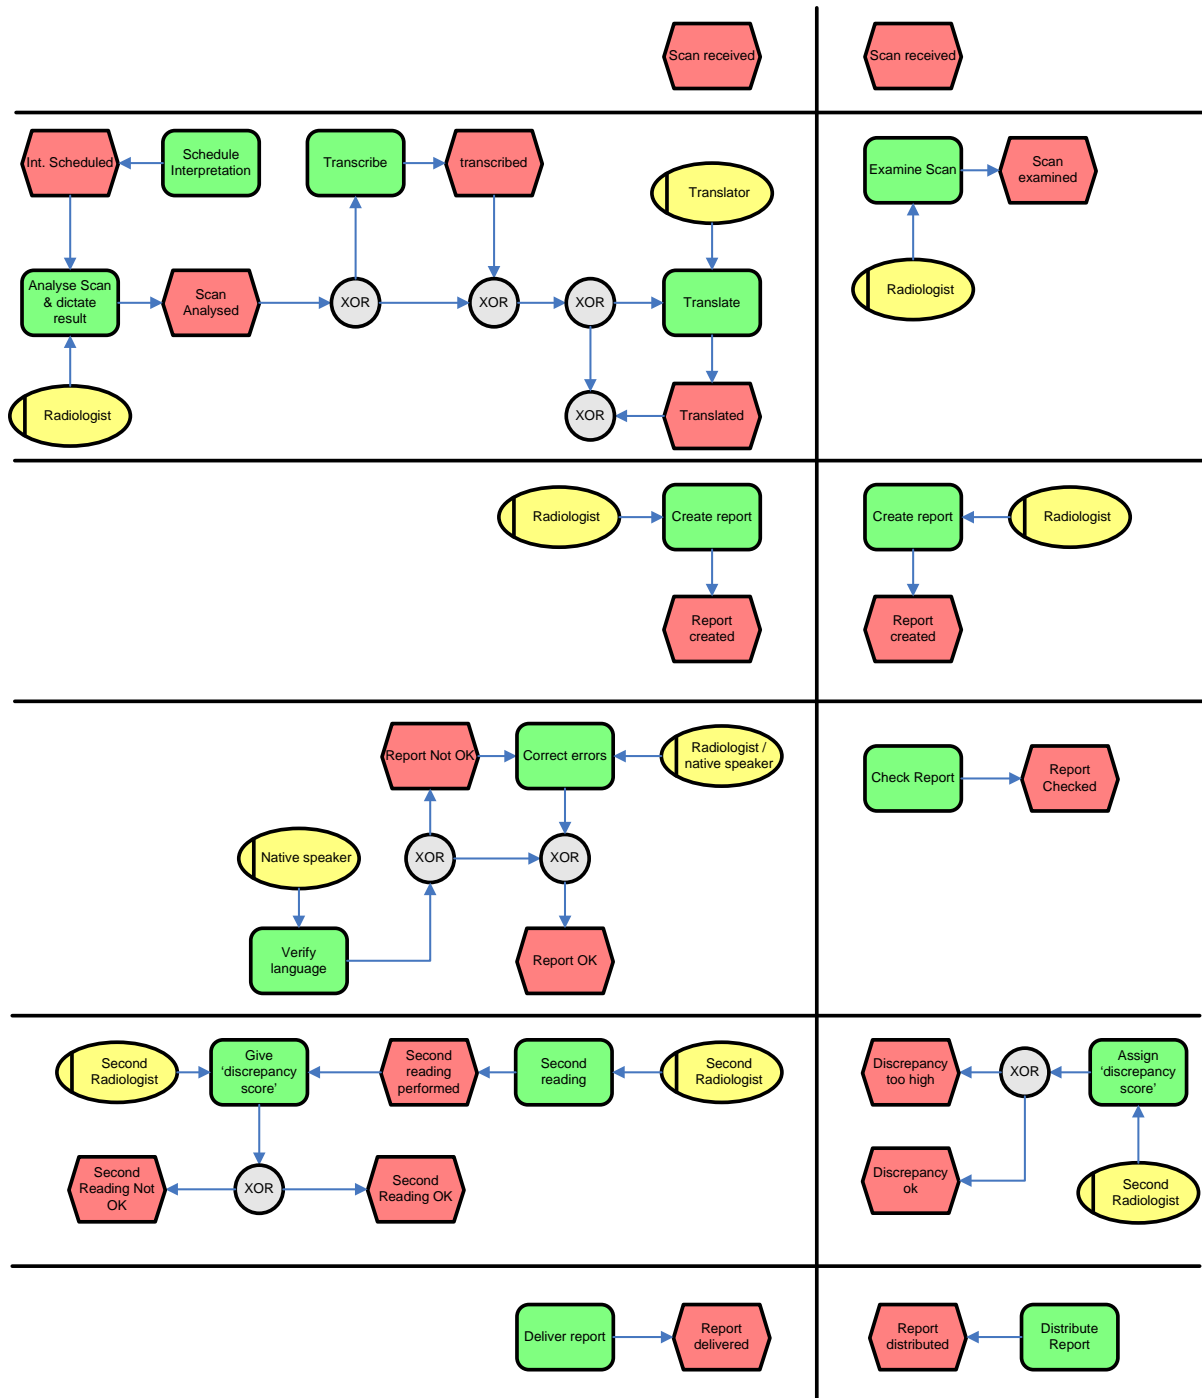


Figure 9: Mapping internal level to external level process

In this case six abstraction relations exist, and they are (see Figure 9):

1. The Scan received event is the same on the internal and external level. This event starts the entire process.
2. The internal level activities *Schedule Interpretation*, *Analyze Scan & dictate result*, *Transcribe*, and *Translate* and the internal level event *Int. Scheduled* are mapped to the external level activity *Examine Scan*. Depending on the internal level activities that are executed (optional choices because of the XOR splits in the model), the external level event *Scan examined* is mapped to from the internal level events *Scan analysed*, *transcribed*, or *Translated*.
3. The internal level activity *Create report* and internal level event *Report created* are mapped one-to-one to the same external level activities.
4. Activities *Verify language* and *Correct errors*, and *Second reading*, together with the event *Report Not OK* are mapped to the *Check Report* activity. Event *Report OK* is mapped to *Report Checked*.
5. The activities *Second reading* and *Give 'discrepancy score'* and the event *Second reading performed* are mapped to the external activity *Assign 'discrepancy score'*. Events *Second Reading Not OK* and *Second Reading OK* are mapped to *Discrepancy too high* and *Discrepancy ok*, respectively.
6. The internal level activity *Deliver report* is mapped to the external level activity *Distribute Report* and the internal level event *Report delivered* is mapped to the external level event *Report distributed*.

The external level process as shown in Figure 8 is one example. Applying different abstractions than the ones shown in Figure 9 leads to a different external level process. For example, the external level activity *Check Report* does not imply that it includes language verification and possible error corrections. This can be agreed upon separately in the e-contract, for example in the service description, but it cannot be monitored as it is not part of the external level process specification. Also, an organization can offer the same service (same external level process) with different values for the QoS aspects in different service offers.

4.3 Established e-Contracts

When an organization needs a service, the requirements are taken by a, so called, matchmaker, which will then search for suitable service offers stored in service repositories. Matchmaking can be performed by the service consumers or service providers themselves, or by an independent third party. A resulting service found by a matchmaker can fall into the following matching categories:

- An exact match. However, as the quality of service specifications can be very diverse, exact matches will be rare.
- An overspecified service offer. A service consumer searches for a service provider, because the provider can perform the service in some ‘better’ way than the consumer itself. Therefore, it is the provider who has the most knowledge and expertise related to the service. Hence the provider can specify more and also more accurately the quality of service attributes corresponding to the service compared to a service consumer that has to specify the service requirements. In this case, the service offer is overspecified compared to the service requirements. Most matchmaker results will fall in this category.
- An overspecified service requirement. Similar to the previous category, however, now the service requirement is more detailed, i.e., it contains more or more precise quality of service attributes, or a more detailed service description. For the same reasons as given in the ‘overspecified service offer’ category, matches will fall rarely in this category.
- Conflicting matches. Although a returned result will match substantially, conflicts in some parts of the service specification exist.

It is then up to the provider to determine whether or not he can meet the requirements, by for example making modifications to the offered service and/or the underlying internal level process. If an exact match is found, an e-contract can be established immediately. However, as the quality of service specifications can be very diverse, exact matches will be rare. Thus, in that case, the matchmaker will return one or more services that are a near match so that negotiations can then take place to see if an agreement can be reached. If an agreement can be reached through negotiations, the e-contract can be established. Figure 10 shows this process of establishing e-contracts.

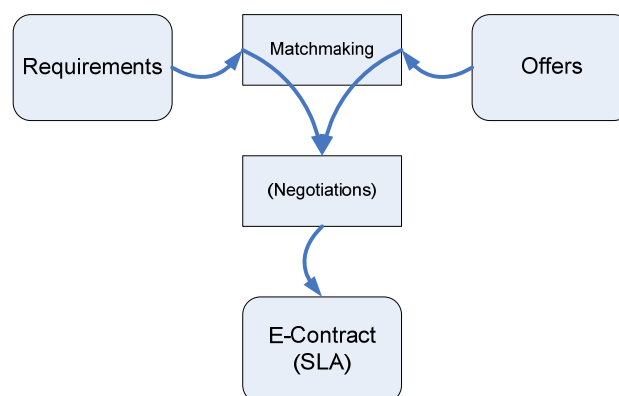


Figure 10: From requirements and offers to e-Contract

For the Teleradiology case, possibly nine different parties are involved (see the service oriented view in Figure 4). All relationships between the parties are bilateral relationships so that the minimum number of contracts required in this business network is eight (i.e., #parties – 1). The type of contracting used, see [GA02], determines the flexibility and dynamism in choosing which parties will be involved in the collaboration. Contracts can be established before or during process execution, also the choice in collaboration partners can be made before or during process execution. If the collaboration parties and the contracts are fixed before process execution starts, then all process executions will proceed in the exact same way. Choosing the collaborating parties at the time that they are required in the process execution, or even establishing the contract at that time increases the flexibility and enables making the best choice at that moment for a specific collaboration partner. As an example, the Clinic or hospital can establish contracts with multiple scan interpreters so it can choose which scan interpreter to use for a specific scan. Also, the scan interpreter can have numerous contracts with different parties that offer the same service (to be able to choose a suitable party when required) or multiple contracts with the same party for different services or for the same service with a different SLA (so that the scan interpreter can choose the collaboration constellation that is required for a particular scan, see also Section 3).

For a collaboration as shown in Figure 4, the following list presents the contracts that are established and indicates what would be included in such contracts. The obligatory standard information in a contract, like party name and address, legal context, etc. is omitted from this list:

1. Patient – Clinic/Hospital (or Hospital – Clinic).

For a patient, this contract contains a layman’s description of the teleradiology process, which can be seen as an external level view on the process. It contains what the patient preparation entails, which scans are required and that a report will be created that contains the an analysis of the scans together. For a hospital (or a medical specialist acting on behalf of the hospital), the contract contains a precise process model (external level view) so that it can be used by the hospital information systems to automatically monitor the progress, synchronize other dependent processes, etc. The process model shown in Figure 2 could be used as an external level view (as it does not contain detailed task specifications for the subprocesses included in this process).

Also the QoS specifications, over the process description/specification, are different for a patient than for a hospital. For a patient, it will include the estimated duration of the relevant part of the process (from prepare patient till the distribution of the report), the chances of having to make additional scans (in case the machine is faulty or the scan failed), the cost of the entire process execution, including insurance reimbursement information, and the possibilities for cancelling the process or rescheduling the scan acquisition.

In a hospital the teleradiology service will be connected to the hospital information systems, which is why all of the QoS aspects listed in Section 4.1 will be included in the e-contract. The establishment of a contract for a teleradiology service invocation can be done ‘at the spot’, i.e., whenever a teleradiology service is required, it can be searched for and the most suitable one can be chosen. The choice can be based on cost, duration, urgency, but also on traveling distance for the patient, etc.

2. Clinic/Hospital – Scan Interpreter

In [Sch06], a currently existing teleradiology setting is presented, in which it becomes clear that the currently existing agreements can be considered a, so called, umbrella contract. Any scan order falls under the agreements specified in the umbrella contract and no specific contract is made for the separate scan orders. The umbrella contract contains:

- a. Pricing, which is determined by the size and length of the contract,
- b. Turn around times (TAT),
- c. Input and output formats,
- d. Second (or also called double) readings, and
- e. Radiologists have to be registered in the country from which the scan originates.

In a service oriented version of the teleradiology process (as described here and in the previous section), contract establishment can be done almost instantaneously and the need for umbrella contracts becomes less. For each required service the most suitable service provider can be searched for and a contract can be established at the time it is necessary. Umbrella contracts can still be established to agree on the service and the QoS for a longer period of time or for a certain (maximum) number of service invocations. Each service invocation will lead to a specific contract instance in which the ‘open’ issues in the umbrella contract have to be filled (also called Partially Filled Contract (PFC) and Completely Filled Contract (CFC) [KGV00]).

3. Scan Interpreter – Transcriber

The scan interpreter has a regular demand for a transcriber service because of the many scan interpretations that need to be transcribed (although speech recognition tools will make this less and less necessary). An e-contract can be established for each service required or umbrella contracts can be made with one or a number of transcribers. The transcribe service is a rather short duration process, of which monitoring is not important. It can therefore be modeled as a black box service. Because of the short duration, level of control is not required. Also, recovery is not required: in case of a process failure the process just needs to be executed again, i.e., the same scan interpretation needs to be transcribed again (and the failed

transcription can be discarded). The actual costs incurred by invoking the service will be handled batch-wise.

4. Scan Interpreter – Translator

The translation service will usually take more time to complete than the transcribe service, because the entire report needs to be translated (compared to transcribing only the dictation of the scan interpretation). Monitoring the progress of this service is therefore desired. Thus, the translation service should provide an external view of their process on which progress monitoring is possible. Because of the duration of this service, a limited level of control is agreed upon. The scan interpreter is allowed to pause or cancel the translation process. Extra costs related to cancelling or pausing the service will have to be paid for by the scan interpreter (specified in the contract). Penalties related to violations of QoS agreements in the contract are also specified in the contract itself. For example, if the duration of a service is longer than the agreed maximum duration, the translator has to compensate the scan interpreter for that (i.e., pay a penalty). This also holds for all other contracts.

5. Scan Interpreter – Native Speaker

A similar reasoning as with the translator above can be applied here and the contract with the QoS specification will therefore be similar to the e-contract with the translator.

6. Scan Interpreter – Second Radiologist

The second radiologist also interprets the scans and then compares his findings with the statements in the report (which implies that the second radiologist must speak the same language in which the report has been written). Depending on the duration of this service and the monitoring requirements, this can be a black box service, with most of the QoS aspects agreed upon in the contract. Unnecessary QoS aspects would be level of control, expected failure rate, and recovery semantics as they are of little use for a short duration service.

7. Clinic/Hospital – Finance Organization

Handling the billing of the teleradiology service is taken care of by another organization (the Finance Organization). The billing service included in the teleradiology process involves creating and sending a bill for the service. The actual payment handling of the bill is done in a separate process within the finance organization. The billing service from the teleradiology viewpoint is therefore a short duration service and can be considered a black box. So, this contract contains the same QoS aspects as the contract between the scan interpreter and the second radiologist.

8. Clinic/Hospital – BI Department

A similar reasoning as with the finance organization above can be applied here and the contract, together with the QoS specification, will therefore be similar to the e-contract with the finance organization.

When an organization collaborates with other organizations (through services), it makes that organization dependent on those other organizations, i.e., the success of a process execution is dependent on the successful process executions of the collaborating organizations. With respect to process reliability, this means that the reliability of process executions is depending on the reliability of the services that are invoked. Thus, it is very important to know the reliability quality of services, so that the most suitable service can be chosen (which is of course a combination of reliability and the other quality of service aspects).

The overall QoS aspects of the process are determined by the separate QoS aspects of the involved services together with the QoS aspects of the consumer service. The more direct collaborations (and thus contracts) exist, the more complex it becomes to determine the overall QoS in case another contract needs to be established. For example, the reliability of the teleradiology process is determined by the reliability qualities specified for the other services involved. If less direct collaborations exist, determining the overall QoS becomes simpler (a party agrees to some QoS which he must ensure even if he outsources as well).

One of the QoS aspects is, as stated before, the reliability of a service execution. The next section introduces the concepts and ideas developed in the XTC project that facilitate the specification and execution support of the reliability QoS aspects of a service, which are subsequently applied in the teleradiology process.

5 Introducing XTC in the Teleradiology process

The main aim of the XTC project is to provide reliability of the execution of contracted services. To achieve this, two key concepts have been developed, which are transactional quality of service and abstract transactional constructs [WGV06, WVG07], see also section 1.2. This section explains these two concepts and illustrates their applicability by describing how they can be applied in the teleradiology process so that the reliability of the process execution is enhanced.

5.1 *Transactional Quality of Service*

Reliability aspects related to service executions are specified in, so called, transactional service level agreements (TxSLA). A TxSLA is a subclass of the regular service level agreement and contains clauses that specify certain reliability requirements (for a service consumer), reliability offers (for a service provider) or reliability agreements (in case an e-contract has been established) related to the service under consideration. A reliability clause (or a reliability agreement) is called a transactional quality of service (TxQoS) specification [WVG07]. Together, the TxQoS clauses in the TxSLA determine the level of process execution reliability for the service to which the SLA refers.

The TxQoS concept consists of the following four attributes, called FIAT [WGV08, WVG07]:

1. **Fluency**: specifies the maximum number of failures in a process execution. If the agreement allows for multiple invocations of the service, fluency can also be used to specify the average number of failures for the total number of invocations allowed by the agreement. The manner in which the occurrence of a failure can be monitored (e.g., while the service is being executed or after the service execution has finished) and the penalty incurred because of the failure are also specified in the agreement (in the TxSLA).
2. **Interference**: specifies what influence a service consumer can exert on the execution of the provider process (e.g., stopping, pausing, and cancelling a process). The timing parameters used to specify when this influence can be exerted, are also specified in the interference TxQoS clauses, as well as the costs related to exerting the influence.
3. **Alternation**: specifies which alternative process paths exist, that can be taken when a failure in a process execution occurs so that the failure can be circumvented. If multiple alternatives are possible, the preference order is also specified.

4. **Transparency:** specifies how much of the provider process can be ‘seen’ by a consumer organization. As explained before, an external level view is agreed upon that is an abstraction of the actually executed internal level provider process (for reasons of secrecy and information overload protection). The transparency attribute provides a measure that indicates the ‘closeness’ of the external level view to the internal level view, i.e., the higher the transparency, the similar the external level process is to the internal level process.

To exemplify the FIAT attributes, they are applied to the teleradiology case. For this, the scan interpretation service is used, but similar FIAT specifications should be included in the TxSLA between the organizations involved in the other service invocations, as shown in the service oriented view of Figure 4. The TxQoS specifications are specified in relation to the external view of the process, shown in Figure 8, and are listed below:

5.1.1 Fluency

Using a statistical model, a so called non-homogeneous poisson process [WGV08], the failure rate of previous executions of the process can be analyzed and the chance that a certain failure occurs can be determined. For the provider organization, this is used to decide what the value of the fluency specification for a service can be, but also how much more to charge for the service if the fluency should be ‘better’ than the average fluency. As fluency indicates the smoothness of the process execution (less failures), it means that a better fluency is specified as a higher value, i.e., the higher the fluency, the smoother the process execution. Therefore, the value of the fluency attribute is specified as:

$$F = 1 / (\text{Expected Failures} + 1)$$

Hence, fluency is specified on a scale from (almost) zero (worst fluency if a high number of failures can occur) to one (highest fluency, if no failures can occur).

For the scan interpretation scenario, suppose that from the statistical analysis it is determined that the chance for one failure (for example, a faulty decompression of the scan image) in the process is 5%, and for two failures (for example, an additional malfunctioning printer that is used to generate the report) the chance is 0.5%. The scan interpreter can now offer this service with a Fluency of 1/2 ($F=1/2$) for a certain price while at the same time offer the same service with a Fluency of 1/3 ($F=1/3$) for a lower price (because it does not have to guarantee a working printer). Additionally, it is also specified what the penalty is in case more failures occur than agreed. For example, this could entail a reimbursement of the service price or a number of free invocations of this service. Monitoring failure occurrences can be done by the consumer organization through runtime monitoring. Any failure will show in the process execution as an deviation from the regular execution flow (regular flow is known

through the external view). Monitoring can also be performed through the analysis of the process execution logs (during or after process execution), in which deviations from the normal execution flows can then be discovered.

For the teleradiology scenario, it is more likely that an umbrella contract is made between the Clinic/Hospital and the scan interpreter in which the fluency is specified in terms of failure chances per batch size of invocations, e.g., a TxQoS specification that states that the maximum number of failures is two per one hundred service invocations ($F=1/3:100$). Even a combination is possible, in which the maximum number of failures for a certain batch size is specified, as well as a maximum number of failures per invocation. For example, five failures is the maximum number of failures that may occur in one hundred service invocations, but not more than one failure may occur in each invocation ($F=1/2;1/6:100$).

Other possibilities allow for a diversification of failures. By having a classification of failures (e.g., critical, severe, mild, non-critical), the service provider can analyze the chance of each failure class and can thus determine what the fluency specifications can be for each failure class in the service offer (including the costs, penalties and monitoring options).

Another possible diversification for the fluency specification is specifying the fluency for each of the activities in the external view (instead of one fluency specification for the entire process). Using the statistical analysis on the past performance of the activities that are part of the service, the chance of failures can be determined for each of them. This can then be stated in the service offer to provide more reliability information to potential consumers.

Fluency only specifies the possible maximum number of failures that could occur in a service execution. However, to be able to determine the impact of failures in certain activities, it is necessary to include, within the TxSLA, information on what the state of the process execution will be after a potential failure has occurred and resolved (after which the process execution can be continued again). An example is given in Figure 11 (an unofficial-extension to the EPC diagram notation).

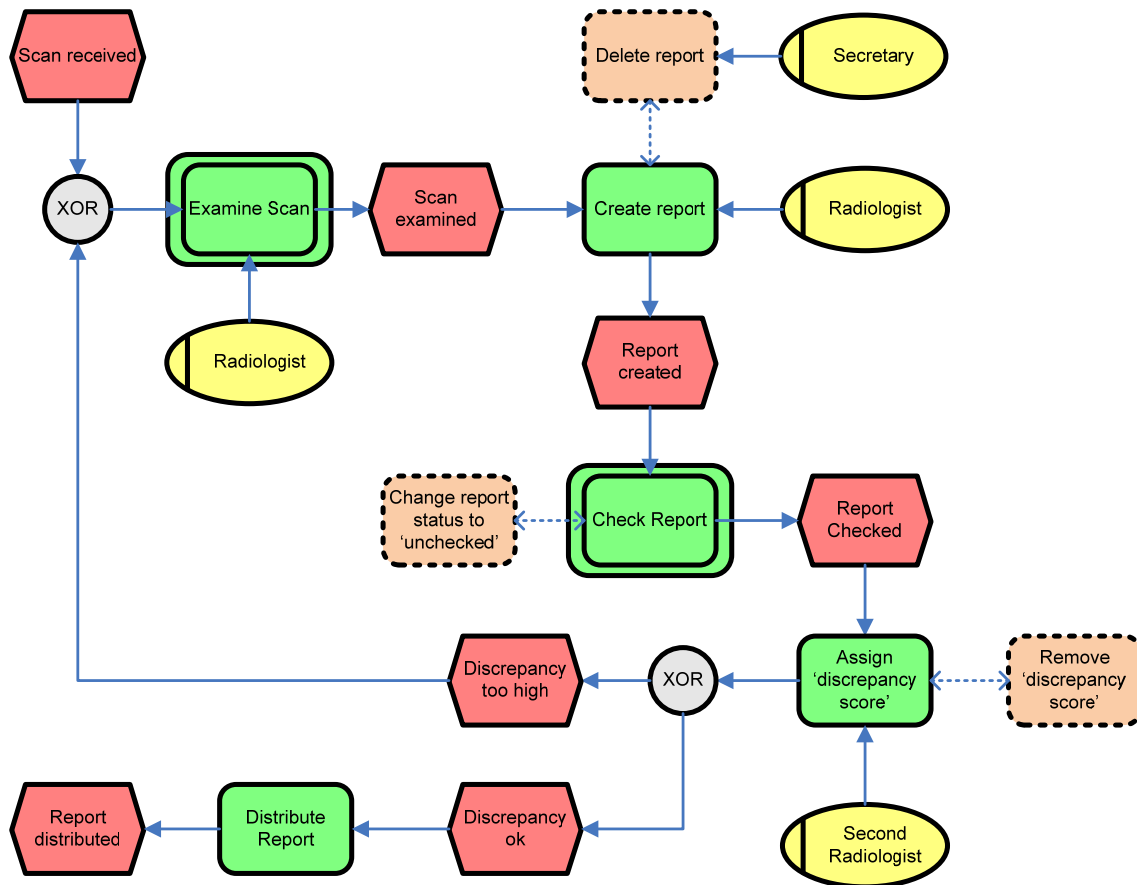


Figure 11: Transactional External View of the Scan Interpretation process

In the figure, the external view of the service is extended with the safepoint and compensation concepts. Safepoints are represented by double lined rounded rectangles and compensation activities are represented by dashed rounded rectangles. Both concepts are introduced to counteract failures or anomalies in the process execution. If a process execution failure occurs, the process execution needs to be 'rolled back' to a process state that is known (pre-specified) to be correct. However, suppose that each activity has already completed and exposed its results. Then these results have to be undone. Undoing activities can be performed through the compensation mechanism. For each activity another activity is specified, called compensating activity, which undoes the work done in the original activity (if possible). By executing these compensating activities in the reverse order of the original activities, the entire process can be undone or 'rolled back'. However, rolling back the entire process is usually not necessary. At some points in the process execution, for which it is known that the process state is correct, the rollback can stop. For this, the safepoint concept is introduced. If a process failure occurs, the process is rolled back to the first suitable safepoint from whereon the normal process execution can start again which could then succeed because of possible changes in the environment so that the failure does not occur again. Note that, because of parallel executing paths, a rollback stops at the first suitable safepoint, which is not necessarily the first encountered safepoint [GVA01].

5.1.2 Interference

The Interference TxQoS attribute is based on the Level of Control concept developed in the CrossFlow project [VG03, GALH01]. As explained before, the interference TxQoS specification can be used to specify the control (or influence) a service consumer can have over the execution of the service at the provider. The possible control consists of pausing a process (i.e., stopping and continuing the process execution), rolling back the process execution, and aborting (or cancelling) the process execution.

Figure 12 shows the external level view of the scan interpretation process, in which the activities are annotated with symbols representing the interference possibilities⁶. The notation used for interference is derived from the well known audio/video control symbols.

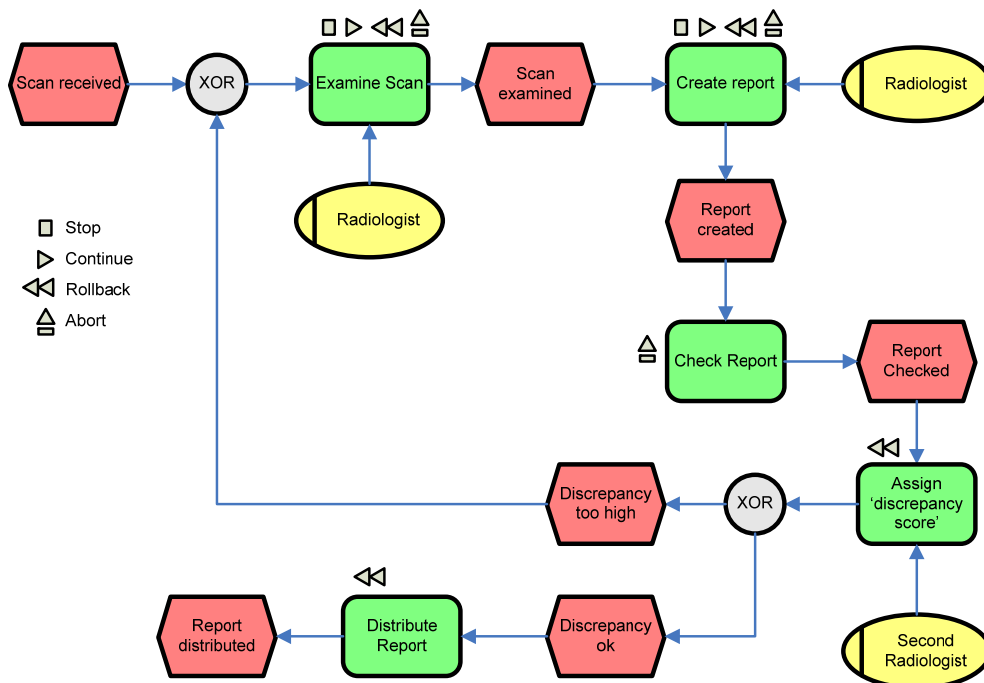


Figure 12: Interference External View of the Scan Interpretation process

Assume that the service shown in Figure 12 is agreed upon by the clinic/hospital and the scan interpreter, i.e., it is included in the e-contract between the two parties. The following examples illustrate the usefulness of the interference TxQoS attribute for the Clinic/Hospital:

- Suppose that, while a scan interpretation is being performed, the status of the patient has changed dramatically. Then the scan interpretation process should be paused so that the requirements for

⁶ As transitions between activities are considered to happen instantaneously, interference is only specified on activities. In this case the states that an activity can be in are ignored. However a more advanced form of interference could take the state of an activity in account. For example, an activity is only allowed to be aborted if it is in the initializing state.

the scan interpretation can be reevaluated in relation to the status change. After pausing the scan interpretation service, it might be continued again, but it might also be cancelled or rolled back. Note that a rollback implies undoing (part of) the process after which it is restarted again. In this example, a rollback can be issued to undo the entire process, after which it can be restarted again, but with the updated status information included (which can change the process execution compared to the first (rolled back) execution).

- While the scan interpretation process is running, the clinic/hospital can abort the process as long as the process execution has not reached the *Assign 'discrepancy score'* activity, thus during *Examine Scan, Create report, and Check Report*. Aborting the process (also called cancellation) means that the process should be undone, but not restarted again. Undoing part of the process, involves executing compensating activities (as explained in the previous section on Fluency). As executing these compensating activities require resources, a price tag is associated with the abort control (in addition to the costs involved for the execution of the process up until the point where the abort is issued). In this case, the abort is not allowed for the last two activities because it costs less to let the process continue than the abort it.

In Figure 12 the interference symbols are used to indicate when a certain control is allowed. The impact of a stop, continue, or rollback can however be on a single activity (or subprocess) or on the entire process. The abort always involves the entire process. For example, if multiple activities are being executed in parallel, then it should be possible to stop one activity, a set of activities, or all of the activities (i.e., the entire process).

Because of the dual role of the scan interpreter in the teleradiology scenario, in which it acts as a provider to the clinic/hospital and consumer to a number of other services, it is important to guarantee consistency between the multiple contracts that such a dual role party is involved in. For example, the checking of the report is outsourced to a native speaker. If in the contract between the scan interpreter and the native speaker it is stated that this service cannot be stopped, then the scan interpreter should make sure that he also does not offer the stop interference to the clinic/hospital (i.e., to his consumers)⁷.

5.1.3 Alternation

To prepare for potential failures that might occur during a service execution that are not correctable, alternative paths in the process can be specified. An example of an uncorrectable failure is the breakdown of a machine. In that case rolling back (part of) the process and restarting will not work as

⁷ The scan interpreter could still offer the stop interference to his consumers, but this would result in a 'stop as soon as possible', as it is not known to the consumer when the native speaker service will finish (the consumer does not even know that the native speaker service exists).

the machine needs to be fixed first. Pausing the process until the machine is fixed might be unacceptable or too expensive to be an option. In such situations, the alternation TxQoS attribute can help. Alternative paths can be specified in the process to circumvent failures. The alternative paths are not the preferred paths through the process, but they might help in successfully completing the process execution. Usually, alternative paths are used in combination with a retriability specification that specifies how many times a specific path in the process execution should be taken (tried) before the failure is considered unrecoverable. For example, in case a failure occurs in a certain activity, then a rollback and restart is tried for a certain number of times. If the retries do not succeed, then an alternative path can be taken.

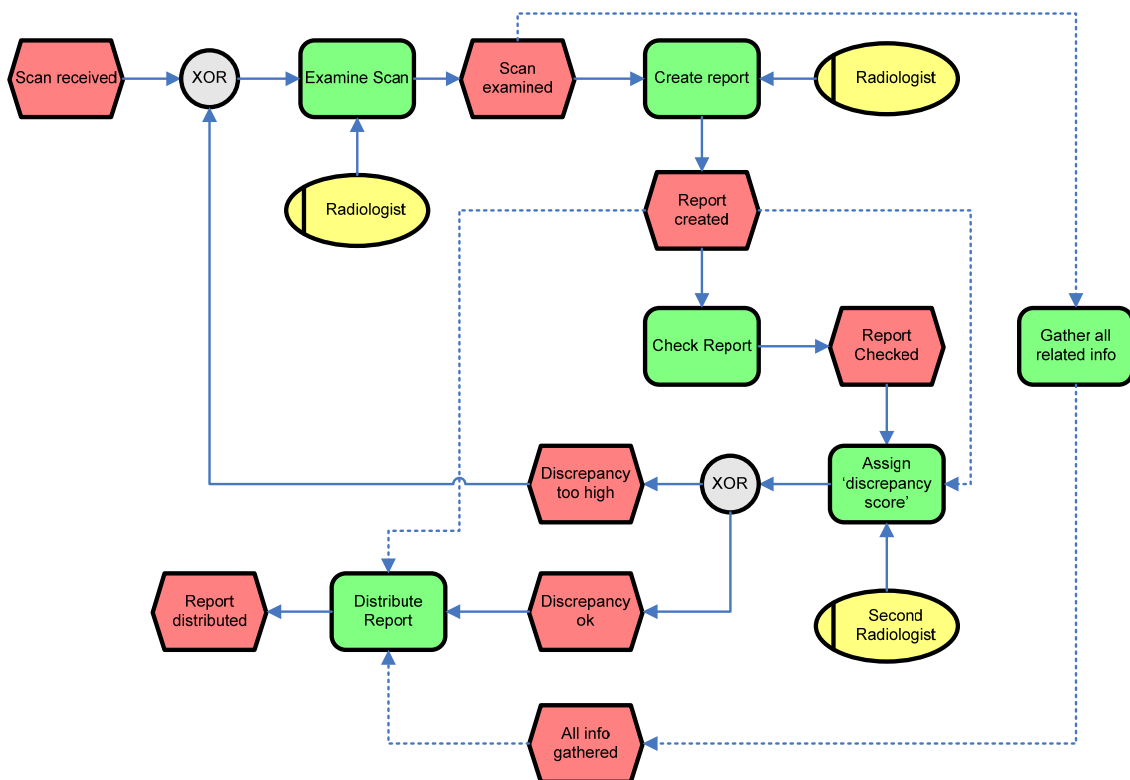


Figure 13: Alternate External View of the Scan Interpretation process

Figure 13 shows the alternatives that are specified for the external view of the scan interpretation service, indicated by the dashed arrows. Suppose that an uncorrectable failure occurs during the *Create report* activity. For example, the radiologist who performed the *Examine scan* activity must be the same person as the radiologist performing the *Create report* activity, but this person is not available for a long time after doing the scan analysis. Then, although it is definitely not the preferred option, the information created during the *Examine scan* activity can be bundled and distributed to the service requestor, via the alternative path that contains the *Gather all related info* activity. In this case, the service requestor gets the scan interpretation information (although not in the desired format and

also not checked) and does not need to cancel this scan interpretation and to search for another scan interpretation service⁸. Also, a discount or other compensation should be made by the service provider as the service has not been executed following the normal specification. Such discounts or compensations are of course specified in the e-contract.

The same line of reasoning can be applied in the situation where the *Check Report* activity cannot complete. In that case two alternative paths can be taken. One alternative skips the *Check Report* activity and continues with the *Assign 'discrepancy score'* activity. The other alternative goes directly to the *Distribute Report* activity, thereby circumventing the *Check Report* and *Assign 'discrepancy score'* activities. Again, the alternatives are not the preferred choice, but at least the process can finish with a result that is satisfactory to the service consumer (the clinic or hospital in this case) as it has been specified in the e-contract.

In the figure, it can be seen that two alternative paths lead from the *Report created* event, which means that a choice must be made to determine which path to take in case alternation is required. Such a choice can be based on the following:

- The location of the failure in the process execution for which the alternative is necessary. For example, in the scan interpretation scenario, if the report has been created and an uncorrectable failure occurs in the assignment of the discrepancy score, then it is only possible to reach a successful end of the process by taking the alternative that skips the checking of the report and the assignment of the discrepancy score.
- Weights that are assigned to the alternative paths. If multiple alternative paths are available to circumvent a certain failure, then the path with the highest weight should be chosen as it has the highest preference. The weights create a preference order among the different alternatives that can be taken from the same point in the process.
- Consumer decision. In the TxSLA can be specified that in case an uncorrectable failure occurs, the service consumer is informed so that he can decide (and instruct the provider) which alternative path to execute. The decision can be based on a number of factors, such as the time that remains till the deadline at which the service must complete, or the costs associated with the paths that are attributed to the service consumer, etc.
- Provider decision. The service provider is the organization that has the most expertise in the activities involved in the service (which is the reason that he can offer a service). Therefore, it is the service provider that can decide on the best alternative path, based on the state that the failed process is in at that point in time.

⁸ Whether an alternative path in the process is acceptable or desirable is domain dependent. In this case, two of the alternative paths circumvent the occurrence of a second reading. This might be acceptable for some of the scan interpretation service consumers, but not for others, in which case these alternative paths cannot be included in the alternation TxQoS specifications.

- Final quality of the end result of the service. As alternative paths are not the preferred path through the process, they have drawbacks compared to the regular path. These drawbacks can influence the quality of the end result. For example, if the checking of the report is skipped, then the final quality of the report is less than if the report had gone through the check and the double reading.
- Always complete as much of the process as possible. In this strategy, the alternative path that leads to a process execution in which as many as possible of the 'regular' activities are completed is chosen.

Similar to the safepoint and compensation concepts introduced for the fluency TxQoS attribute in Section 5.1.1, a mapping is required that maps the external level alternative paths to the internal level alternative paths (see Section 5.3).

5.1.4 Transparency

The transparency TxQoS attribute represents a measure of abstraction, i.e., it indicates how much of the internal level process is exposed in the external level process. The following (non-exhaustive) list presents the most common ways in which the transparency TxQoS can be specified:

- As a measure between the number of internal activities and external activities. For the scan interpretation process, the number of internal level activities is ten and the number of external level activities is five. The transparency is thus 50% (5 divided by 10).
- Similar to the ratio between internal and external level activities, the number of events on both levels can be taken for the transparency measure. For the scan interpretation process the transparency would then be: 7 (external level events) divided by 12 (internal level events), which is around 58%
- Dividing the number of internal control nodes through the number of external control nodes is another measure for transparency. The transparency for the scan interpretation process then is around 20% (2 divided by 10)
- The number of transitions (arrows in the models) can also be used as a transparency measure. The internal level scan interpretation process has 37 transitions and the number of transitions on the external level is 14. Transparency based on the transition ratio is therefore 38%.
- Naturally, any combination of the four possibilities above also leads to a measure for the transparency. For example, the average transparency for the activities, events, and control nodes is 43% $((0.5 + 0.58 + 0.2) / 3)$, or the transparency of the service over the activities and events is 55% $((5+7) / (10+12))$.

- Analogous to the transparency specified on the basis of control flow elements (activities, events, transitions), data elements can also serve as a basis for transparency. This can then in turn also be combined in the other transparency measures based on control flow or possibly other process perspective elements (for example resources, or organizational units).
- A further differentiation of the elements used in the transparency calculation may lead to a more suitable measure. For example, taking the duration of the activities into account, or give more weight to activities that are a direct (i.e., a one-to-one) mapping from the internal level to the external level. In the scan interpretation service, the *Create Report* activity is an activity that exists on both the internal as well as on the external level (a one-to-one mapped activity) which can be considered to increase the transparency more than just counting the number of activities.

As can be seen or derived from the examples given in the list above, the transparency TxQoS attribute can be specified for the entire service. However, transparency can also be specified for the activities, events, or data elements separately. For example, the transparency for the *Examine Scan* activity is 25% (4 internal level activities are abstracted into 1 external level activity), or the transparency of the *Create Report* activity is 100%.

The FIAT attributes are used to specify reliability aspects on the service level, i.e., using business terminology. However, the internal level process is the process that is actually executed, the reliability of which is governed by the transaction management support associated with that internal level process. Transaction management support for a process is specified in technical terminology which has been captured in the abstract transactional construct concept that is explained in the next subsection.

5.2 Abstract Transactional Constructs

TxQoS increases the reliability by facilitating explicit agreements related to reliability (using the FIAT attributes) in service level agreement(s). The reliability of actual process executions is increased through the support of transaction management. Transaction management has already been used for decades to provide reliability in databases. A transaction is a number of operations (activities or state changes) grouped together with an explicit start and an explicit end. In case an exception (failure) occurs within a transaction, the outcome (or behavior) of that transaction is determined by the transaction model. For example, the transaction model might dictate that all operations in the failing transactions should be undone and then retried. Advances in transaction management make it also applicable in the area of process management, thereby enhancing the reliability in process executions

through the support of advanced transaction models. The most well-known and useful transaction concepts are:

- **Atomicity:** All operations within an atomic transaction must be done or none at all.
For example, in the scan interpretation process the activities *Second reading* and *Give 'discrepancy score'* must be completed together (or none of them). In other words, if the assignment of the discrepancy score fails, then also the second reading fails and must be redone. Both activities are therefore assigned to one atomic transaction.
- **Isolation:** The intermediate results from the operations performed within a transaction can not be seen by other transactions.
Besides being part of one atomic transaction, the transaction containing the activities *Second reading* and *Give 'discrepancy score'* should be isolated. This prevents that other concurrently executing process (activities) can see the intermediate results of those two activities, but can only see the result after the discrepancy score has been assigned and that activity (and the transaction) has finished.
- **Compensation:** Semantically undoing the result of the operations within a transaction.
Similar to the explanation of compensation in Section 5.1.1, the activities on the internal level process can also have compensating activities associated with them. Figure 14 shows the compensating activities as they are assigned to the 'regular' activities of the main teleradiology process. Note that the model has been simplified, compared to the model shown in Figure 2, for clarifications (the events that are left out are not relevant for compensations). As can be seen, for the processing of the scan (the *Process Scan* activity that represents the scan interpretation service) a compensating activity has been assigned as well. This compensating activity can be used to undo the entire outsourced interpretation service by the clinic/hospital itself. Also a subprocess has been added that contains the preparation of the patient and the acquisition of the different scans, which can be compensated in its entirety by one compensating activity: *Archive scans & Update status*.
- **Retriability:** If a transaction fails, retry its execution a specified number of times, so that the chance of a successful termination of the transaction is increased. If the transaction still fails after the maximum number of retries, then abort the transaction (i.e., the transaction fails).
Suppose that the clinic owns two CT scanners and that one of the scanners fails. This would cause the processes, in which the failing scanner is scheduled to be used, to fail as well. However, suppose that it is a random defect in the scanner that occurs very infrequently (and repairs are due later), then the scan acquisition can be retried (as the

patient is prepared to undergo the acquisition anyway), which might result in a successful scan acquisition after all.

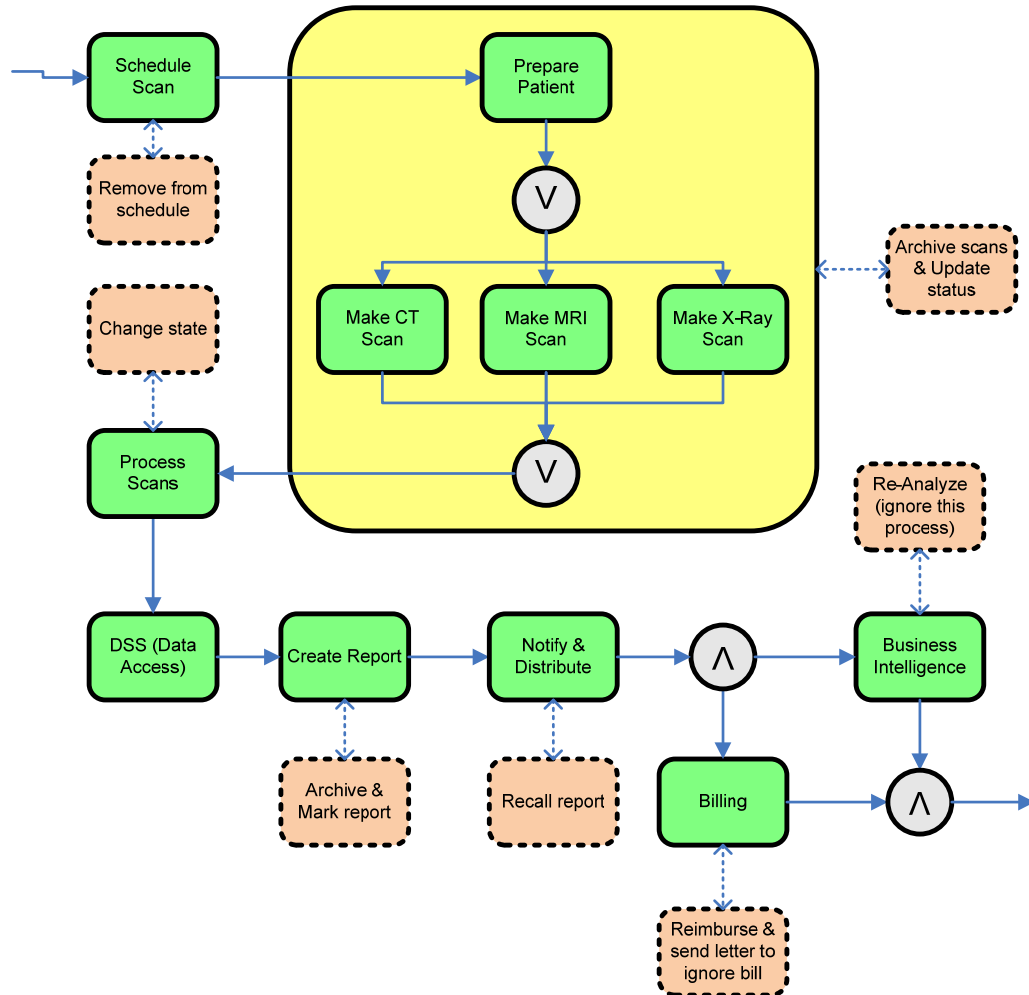


Figure 14: Compensations for the Main Teleradiology process

- Alternations: If a transaction fails, alternative operations can be performed (alternative in relation to the failing operations) so that the transaction might terminate successfully after all.

Continuing on the retriability example above, suppose that the scan acquisition fails again and again (retrying), then after a certain number of retries, an alternative can be taken, which consists of rescheduling the scanners so that the failing scanner is unscheduled and the second (backup) scanner is used in the schedule. Then the scan acquisition can terminate successfully and the scan is acquired, although the original scheduled scanner failed. Another alternative could have been to skip the acquisition of this scan, but that might result in a diagnosis that is less reliable, because not all required scans (indicated in

the scan order) have been completed and thus less information to base to diagnosis on is available.

- Nesting⁹: Nesting operations in a transaction introduces a hierarchical structure in the transaction. Child operations can be executed in parallel and the intermediate results are visible to other branches in the nested transaction (depending on the visibility specifications of the nested transaction model used). Parent operations decide on the success or failure of its child operations.

The scan acquisition subprocess, shown in Figure 14, in which the patient is prepared and the scans are acquired can be supported by a nested transaction in which the acquisition of the scans can be done in parallel (almost parallel as the patient must be physically present, however the pre- and post-processing of the scanners can be done in parallel). This way, intermediate results can be seen within the different activities (staff at the MRI Scanner can see whether the X-Ray scan has been almost completed so that they can expect the patient within a small period of time). Also, the scan acquisition subprocess can still finish successfully, even if not all scans have been acquired. See also the example given for the alternative transaction concept.

Numerous transaction models have been developed, each with their own transaction semantics by employing some of the transaction concepts explained above. However, the ATC concept has been developed to abstract away from the specifics of certain transaction models, by just looking for the common and important transaction concepts (or mechanisms) that are useful to support the execution of business processes. So, by assigning ATCs to (parts of) the process (i.e., using transaction concepts, as the ones explained above), the reliability of the teleradiology process can be increased. It is known beforehand (specified) what the effects on the process execution will be in case a failure in the process occurs. Assigning ATCs to the process leads to a transaction tree that determines the transaction behavior of the process. For the teleradiology, an example transaction tree (or ATC tree) is given in Figure 15 below. An ATC is represented by a rectangle, denoting the transactional semantics or transaction model required for the specific activity (stated between parentheses). Each activity or (sub)process of the teleradiology process is supported by a specific ATC that support that activity with the required transactional semantics. The leaf-nodes in the Tx-Tree correspond to the activities in the teleradiology process, as shown in Figure 2. Intermediate nodes correspond to subprocesses (which are not shown in the process model of Figure 2) and the top ATC represent the teleradiology process itself. Details of specific transaction models are out of scope for this report, but can, for

⁹ Although nesting is an entirely different class of transaction concept compared to the concepts explained before, it is explained here as it is a very useful transaction concept to support business processes.

example, be found in [WV+08, NR+07, GR93, GS87, Moss81]. As can be seen in the Tx-Tree, it is also possible to have no transactional support for an activity. This is indicated by the ‘Empty’ ATC. As the Tx-Tree corresponds to a process model, the order between the ATCs is omitted as it follows from the order of the activities in the transaction model and would therefore be redundant.

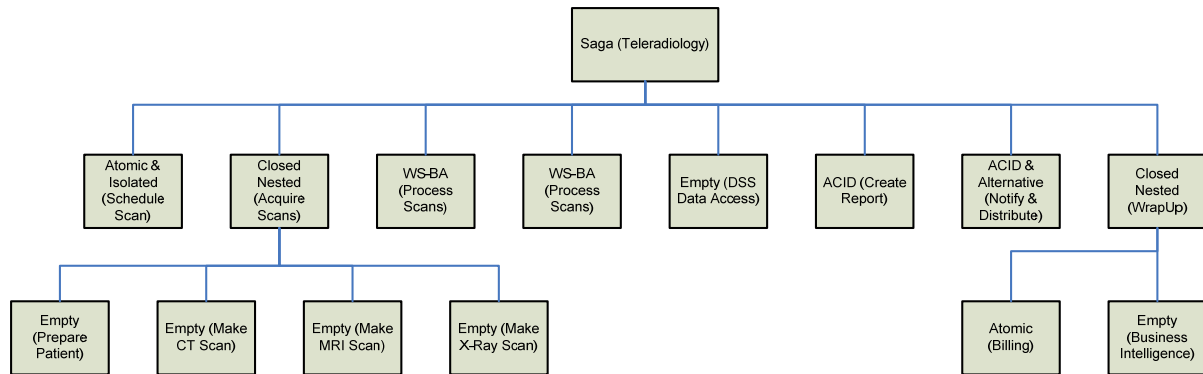


Figure 15: Teleradiology ATC Tree

From the reliability on the business level, i.e., the FIAT attributes, and the reliability on the internal level, i.e., the ATCs supporting the internal level process, it is clear that a mapping between the two must exist that ensures the consistency between the reliability specified on the two different levels.

5.3 Mapping

Before an organization can offer a service to other parties, the internal level process needs to be mapped to an external level process view. However, besides the mapping of the process itself, other aspects associated to the process, like reliability aspects, need to be mapped as well (if relevant for the service). Consistency between the two levels should be guaranteed so that the internal level process execution can assure the agreements made on the external level and that the external level represents the status and execution progress as it actually is on the internal level. The mapping between the two levels is a bidirectional mapping. This means that it should be able to map consumer requirements of the service (or specific provider wishes for the service) to the internal process level, possibly modifying the internal level to match the requirements. Vice versa, it should be able to map the internal process level to the external level, whereby the QoS associated with the service is determined. Note that details on this mapping are out of scope for this report.

Figure 16 shows the internal level scan interpretation process (as already shown in Figure 3), complemented with compensating activities and one alternative path (leaving out the roles involved in the process execution, as it is not relevant here).

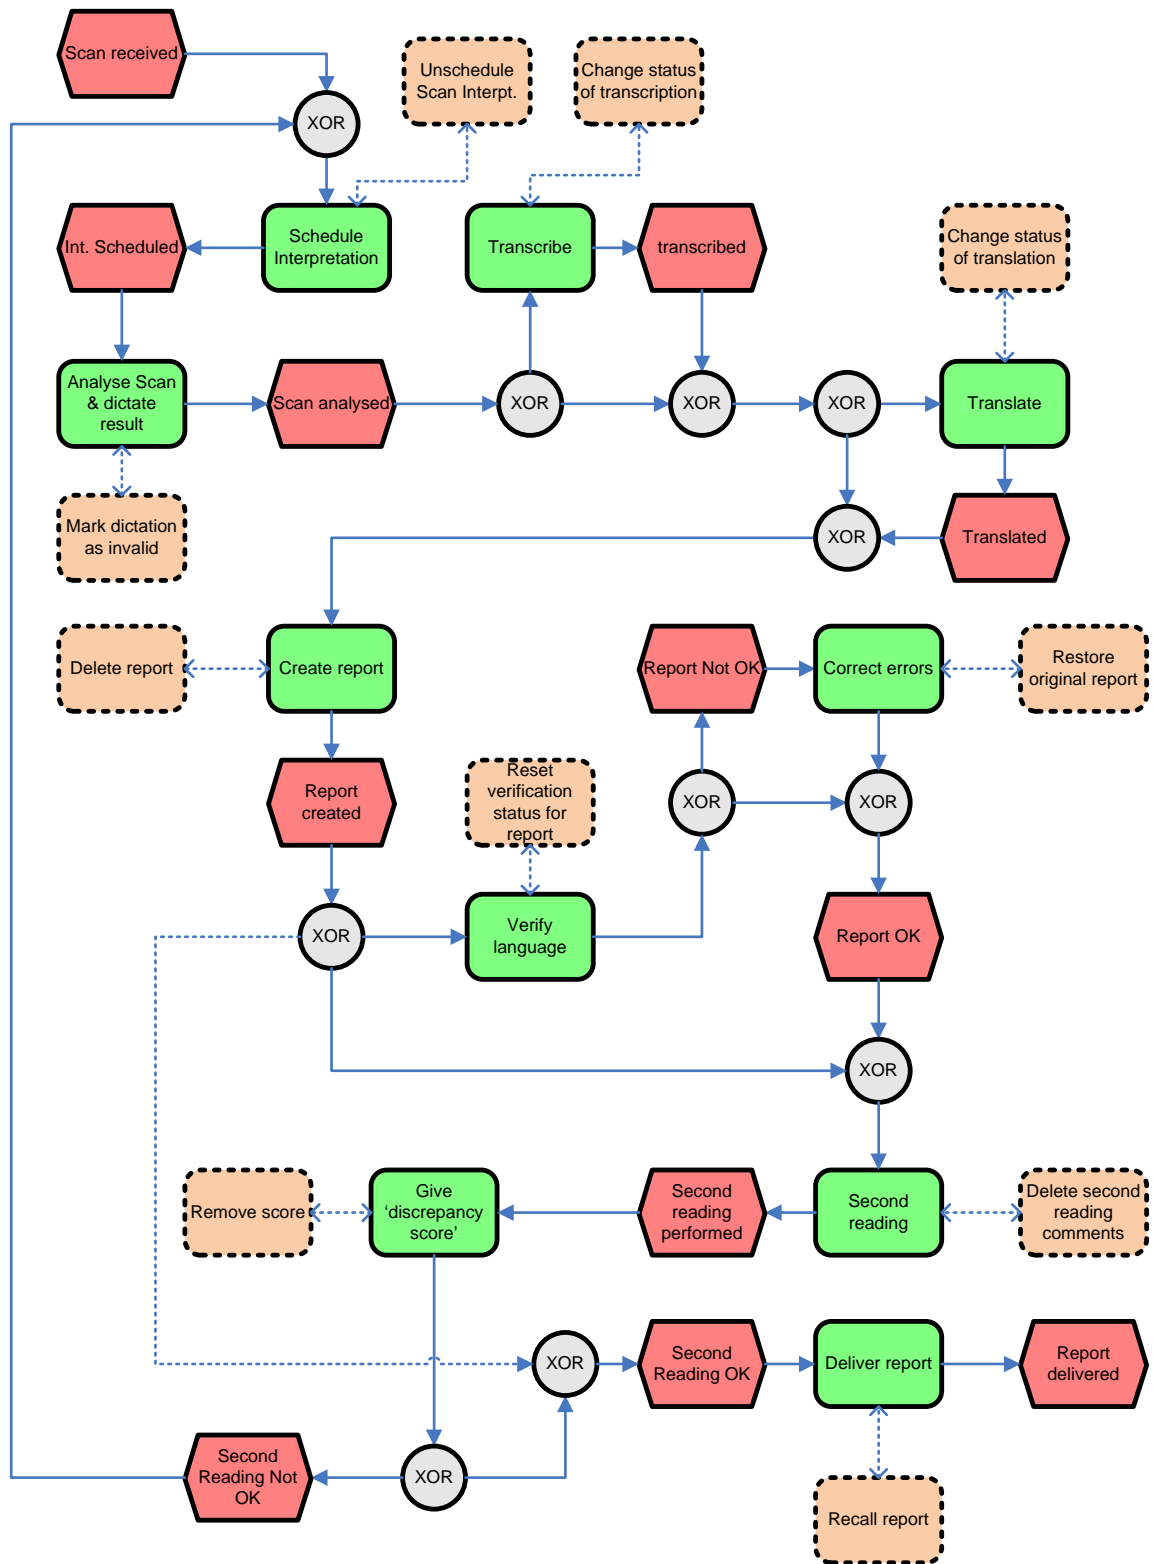


Figure 16: Internal Scan Interpretation process with Compensation and Alternative

Suppose the consumer requirements (clinic or hospital) on a scan interpretation service are represented by the figures shown in Section 5.1. Thus, the consumer has specific requirements related

to compensations, level of control, and alternatives. Similar to the mapping between the external and internal level process, shown in Figure 9, a mapping exists between the external and internal level compensating activities. This mapping is shown in Figure 17. From the figure it can be seen that not all internal level compensating activities are encapsulated by an external level compensating activity. This means that if a failure occurs in the internal level, then this will not be exposed to the external level, in other words, it is hidden to the service consumer.

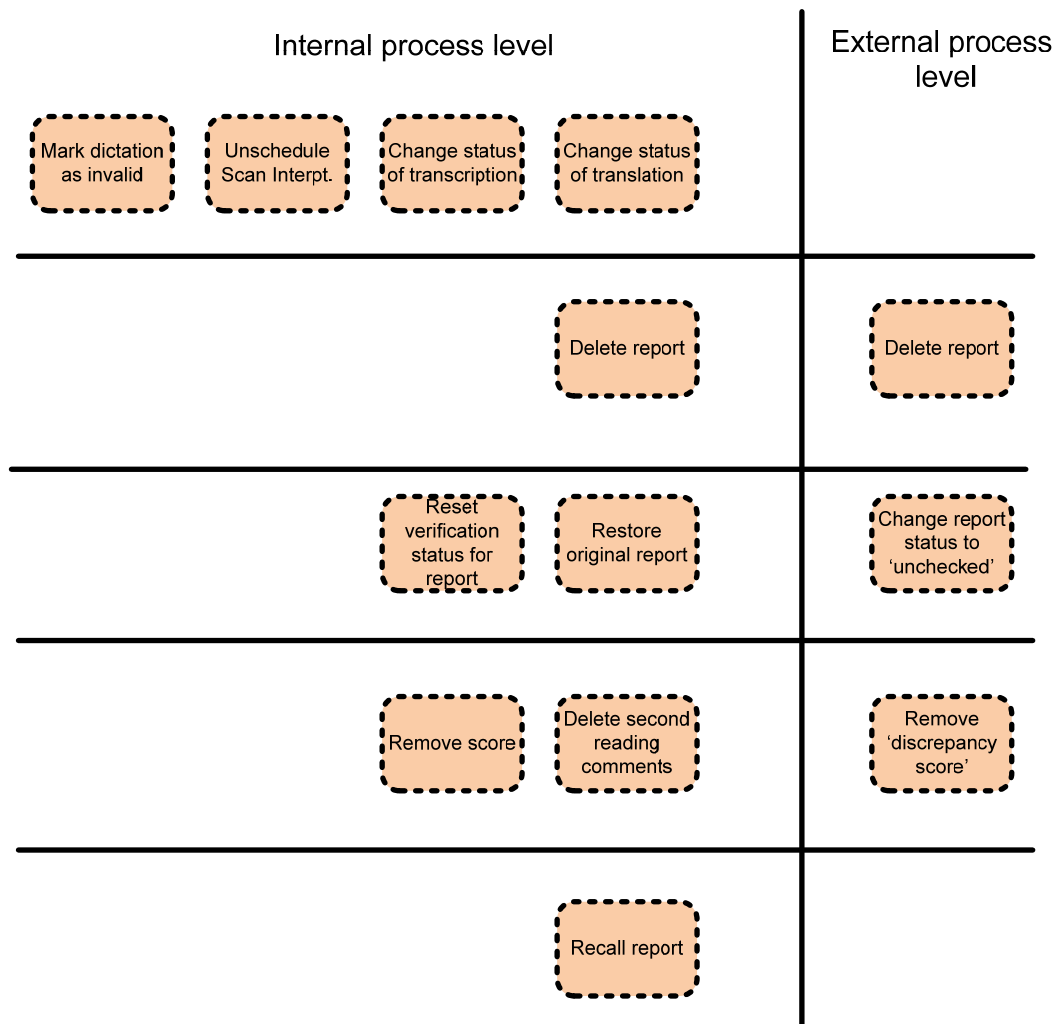


Figure 17: Internal compensation versus external compensation

The external level view of the scan interpretation process contains two safepoints (*Examine Scan* and *Check Report*), i.e., the service consumer requires the two safepoints as indicated in Figure 11. However, as the internal level process is not supported by safepoints the external level safepoints cannot be 'offered' without modifying the internal level process by placing (at the correct locations) safepoints in the internal level process specification to support these external level safepoints. Note

that the designer of the internal level process omitted the internal level safepoints for a reason, so adding them to match the service consumer requirements might not be possible, if it cannot be justified against the reason to omit them in the first place¹⁰. For example, the Examine Scan activity is a safepoint, so that the ‘ending’ activities that it encapsulates must also be ‘marked’ as a safepoint. In this case these activities are: *Analyse Scan & dictate result*, *Transcribe*, and *Translate*. For the *Check Report* safepoint activity, the internal level activities *Verify language* and *Correct errors* must be ‘marked’ as safepoints. This ensures that a rollback on the external level (semantically undoing what has been done) that stops at a safepoint is consistent in case the rollback itself is being executed on the internal level (using the internal level compensating activities).

The external level view contains three alternate TxQoS specifications. However, only one alternative is specified on the internal level. This requires a modification of the internal level specification (if a match to the consumer requirements is desired) in which the other two alternatives must be added. What these alternatives look like is to be designed by the service provider (scan interpreter), as long as it reflects the required alternatives of the external level. The already indicated internal alternative path corresponds to the external alternative path that skips over the checking and second reading of the report. The external alternative path that only skips the checking of the report and continues with the second reading corresponds to a normal process transition in the internal level process. This transition should thus be changed in an internal level alternative path. Note that apparently an inconsistency existed between the internal and external level process models with respect to this transition; if this transition would have been executed it would have reflected on the external level as skipping the *Check Report* activity, which is impossible according to the external level process specification. The third alternative path can just be added to the internal process specification, in which case also one or more new activities need to be added that are abstracted into the external *Gather all related info* activity.

¹⁰ In the usual case, it is the service provider who maps the internal level transaction concepts, like safepoint, to the external level, but to match specific requirements the reverse can sometimes also occur.

6 Summary and Conclusions

In this report, a case study conducted within the teleradiology domain has been described. Starting with a process model that describes the current state-of-the-art way of working in the teleradiology domain, this model has been adapted, modified, and improved to illustrate a possible future scenario for the teleradiology domain. First, the process models were analyzed using a service oriented view perspective. This allowed the identification of services involved in the teleradiology process. Through the identified services, the service interactions between involved parties become explicit facilitates the identification of possible collaboration constellation (with the associated advantages and disadvantages). As the service oriented view explicitly indicates where collaborations between parties (can) exist, it also indicates where agreements between these parties are required. Agreements are captured in electronic contracts (or e-contracts) which are established by matching the service requirements posed by a service consumer with the service offers made by service providers. Service contracts, i.e., e-contracts concerning services, contain service level agreements (SLA) that govern the quality aspects of a service. The quality aspects covered in this report deal with service execution reliability, as designed in the XTC project. The transactional quality of service attributes (TxQoS) defined in the FIAT framework are used to specify service reliability aspects that are agreed upon by the service consumer and provider and are hence included in the e-contract (within the transactional service level agreement, or TxSLA). The internal level provider process is supported by ATCs which provide the required transactional behavior of that specific part of the process that they are assigned to. The mapping between the FIAT attributes and the ATCs ensures that the agreed upon reliability is guaranteed over the internal level provider process (and vice versa). The application of the introduced XTC concepts within the teleradiology process models has been incrementally illustrated. Appendix A presents the teleradiology process models in which the XTC concepts are combined into one model.

This report makes clear that a number of improvements are feasible within the healthcare domain, through the application of the ideas and concepts presented here. Services, and especially their corresponding quality aspects, are becoming more and more important in the healthcare domain. The service oriented view introduced in this report forces an explicit identification of services that are required (or already used in an existing situation). This facilitates choosing collaborating parties and establishing e-contracts at the time they are required, thereby greatly increase the flexibility in composing end-to-end healthcare processes. The service level agreements contained in the e-contracts capture the agreed quality aspects related to the service. The quality aspect focused on in this report concerns service execution reliability, which can be improved (as it is currently hardly addressed) and guaranteed through the FIAT framework and the ATC support. However, other quality of service aspects can be covered through analogous adaptation and modification of the healthcare process.

Acknowledgements

Thanks got to Professor Krishnamurthy Vidyasankar for his comments and feedback on a draft version of this report.

7 References

- [AG04] S. Angelov, P. Grefen; *The Business Case for B2B E-Contracting*; Proceedings 6th International Conference on Electronic Commerce, ICEC'04: Towards a New Services Landscape; Delft, The Netherlands; 25-27 October, 2004; pp. 31-40
- [AG03] S. Angelov, P. Grefen; *The 4W framework for B2B e-contracting*; International Journal Networking and Virtual Organisation, Vol. 2, No.1, 2003; Inderscience; pp.78-97
- [GA02] P. Grefen, S. Angelov; *On τ -, μ -, π -, and ε -Contracting*; Proceedings CAiSE Workshop on Web Services, e-Business, and the Semantic Web; Toronto, Canada; 27-28 May, 2002.
- [GALH01] P. Grefen, K. Aberer, H. Ludwig, Y. Hoffner; *CrossFlow: Cross-Organizational Workflow Management for Service Outsourcing in Dynamic Virtual Enterprises*; Data Engineering Bulletin; Vol. 24, No. 1; IEEE; 2001; pp. 52-57.
- [GLA03] P. Grefen, H. Ludwig, S. Angelov; *A Three-Level Framework for Process and Data Management of Complex E-Services*; Cooperative Information Systems; Vol. 12, No. 4; World Scientific; 2003; pp. 487-531.
- [GMK07] P. Grefen, N. Mehandjiev, G. Kouvas, G. Weichhart, and R. Eshuis; *Dynamic Business Network Process Management in Instant Virtual Enterprises*; BETA Technical Report Series WP-198, Eindhoven University of Technology, 2007.
- [GR93] J. Gray, A. Reuter; *Transaction Processing: Concepts and Techniques*; Morgan Kaufmann, San Mateo, 1993.
- [GS87] H. Garcia-Molina, K. Salem; *Sagas*; Proceedings of the SIGMOD Conference, ACM, New York, 1987; pp. 249–259.
- [GVA01] P. Grefen, J. Vonk, P. Apers; *Global Transaction Support for Workflow Management Systems: from Formal Specification to Practical Implementation*; VLDB Journal; Vol. 10, No. 4; Springer; 2001; pp. 316-333.
- [KGV00] M. Koetsier, P. Grefen, J. Vonk; *Contracts for Cross-Organizational Workflow Management*; Proceedings 1st International Conference on Electronic Commerce and Web Technologies (EC-Web); London, UK; 2000; pp. 110-121.
- [KNS91] G. Keller, M. Nüttgens, A.-W. Scheer. *Semantische Prozessmodellierung auf der Grundlage „Ereignisgesteuerte Prozessketten (EPK)“*; Veröffentlichung des Institut für Wirtschaftsinformatik, Paper 089, Saarbrücken, 1991 <http://www.iwi.uni-sb.de/iwi-hefte/heft089.ps>
- [Moss81] J.E.B. Moss; *Nested transactions: an approach to reliable distributed computing*; PhD thesis, MIT, April 1981. Also available as Technical Report MIT/LCS/TR-260.

- [NR+07] E. Newcomer, I. Robinson, T. Freund, M. Little; *Web services business activity framework (WS-BusinessActivity) version 1.1*; OASIS, April 2007. <http://docs.oasis-open.org/ws-tx/wstx-wsba-1.1-spec-errata-os.pdf>
- [Nua08] Nuance Communications, Inc.; Dictaphone Healthcare Solutions; 2008. <http://www.nuance.com/dictaphone/>
- [PIZ08] Inspectie voor de Gezondheidszorg, Vereniging van ziekenhuizen, Nederlands Federatie van Universitair Medische Centra, Orde van Medisch Specialisten; *Prestatie-indicatoren ziekenhuizen, Basisset 2008*. <http://www.snellerbeter.nl/prestatie-indicatoren/>
- [Pre08] PreScan B.V.; PreScan Medische Check-Up MRI/CT Scans: Preventief Medisch Onderzoek; 2008. <http://www.prescan.nl/SiteFiles/PDF/Prescan%20Brochure.pdf>
- [Ran04] E.R. Ranschaert; *Advancing European TeleRadiology*; European Hospital Volume 13, Number 2, 2004.
- [Roy08] Royal Hospitals website, 2008. <http://www.royal-hospitals.com/>
- [Sch00] A.W. Scheer; *ARIS Business Process Modeling*; Springer Verlag, 2000.
- [Sch06] Jan Schillebeeckx; *Eurad Consult pioniert met teleradiologie*; MemoRad 11(3), 2006. <http://www.radiologen.nl/43/memorad>
- [Swe07] Marcel Swennenhuis; De veranderende processen en rollen in radiology t.g.v. digitalisering; Beta Seminar IT in Healthcare operations management, October 2007. http://fp.tm.tue.nl/beta/events/healthcare%20041007/index_workshop.htm
- [Udd08] Universal description, discovery and integration (UDDI), 2008; <http://uddi.xml.org>
- [VG03] Jochem Vonk, Paul Grefen; *Cross-Organizational Transaction Support for E-Services in Virtual Enterprises*; Distributed and Parallel Databases; Vol. 14, No. 2; Kluwer Academic Publishers; 2003; pp. 137-172.
- [VW+08] Jochem Vonk, Ting Wang, Paul Grefen, Floor Haak – van der Lely, Jan van Aarle, and Suzan Brugmans; *An analysis of contractual and transactional aspects of a cardiothoracic surgery process*; Beta Technical Report Series WP-244, Eindhoven University of Technology, June 2008.
- [VWG07] Jochem Vonk, Ting Wang, and Paul Grefen; *A Dual view to Facilitate Transactional QoS*; in Procs. of the 16th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative enterprises (WETICE 2007); June 18-20, 2007, Paris, France, 2007
- [WGV06] Ting Wang, Paul Grefen, Jochem Vonk; *Abstract Transaction Construct: Building a Transaction Framework for Contract-Driven, Service-Oriented Business Processes*; In Proceedings of the 4th International Conference on Service Oriented Computing (ICSOC). Amsterdam, the Netherlands, 2006.

- [WGV08] Ting Wang, Paul Grefen, Jochem Vonk; *Ensuring Transactional Reliability by E-Contracting*; In 20th International Conference on Advanced Information Systems Engineering (CAiSE); Montpellier, France; Vol. LNCS 5074; Springer, 2008; pp. 262-265.
- [WVG07] Ting Wang, Jochem Vonk, Paul Grefen; *TxQoS: A Contractual Approach for Transaction Management*; In Proceedings of the 11th IEEE International EDOC Enterprise Computing Conference (EDOC2007). Annapolis, MD, USA. IEEE Computer Society, 2007; pp. 327-338.
- [WV+08] Ting Wang, Jochem Vonk, Benedikt Kratz, Paul Grefen; *A survey on the history of transaction management: from flat to grid transactions*; Distributed and Parallel Databases, Vol. 23, Springer, 2008; pp. 235-270.

Appendix A Summary of the processes

In this appendix a summary is presented of the scan interpretation in which the application of the XTC concepts and ideas are combined and main teleradiology processes that includes the view on the outsourced scan interpretation service.

Appendix A.1 The Scan Interpretation Process

Although most of the focus in this report was on the scan interpretation and main teleradiology processes, it is equally applicable for the other services and parties that are involved in the entire teleradiology service chain.

In summary, the organization, that offers the scan interpretation service, creates an external view of its process, which it will offer as a service to potential service consumers. A consumer organization can use this external view to monitor the service execution and follow the progress. Figure 18, shows the external view of the scan interpretation process as it is offered. It contains the process itself, but also compensating activities (to indicate failure resolution to consumers), alternative paths that are possible during the process execution, and the level of control that it offered to consumers. Both the alternative paths and the level of control are fully specified using the alternate and interference transactional quality of service specifications. Together with fluency and transparency specifications, they form the FIAT TxQoS attributes, which is part of the transactional service level agreement (TxSLA).

Note that the external level view is an abstraction of the internal process (the process that is actually executed), the service provider can therefore create and offer different external level views of the same internal process, so that the requirements of different consumers can be satisfied or to assign different values for the quality of service specifications to the services.

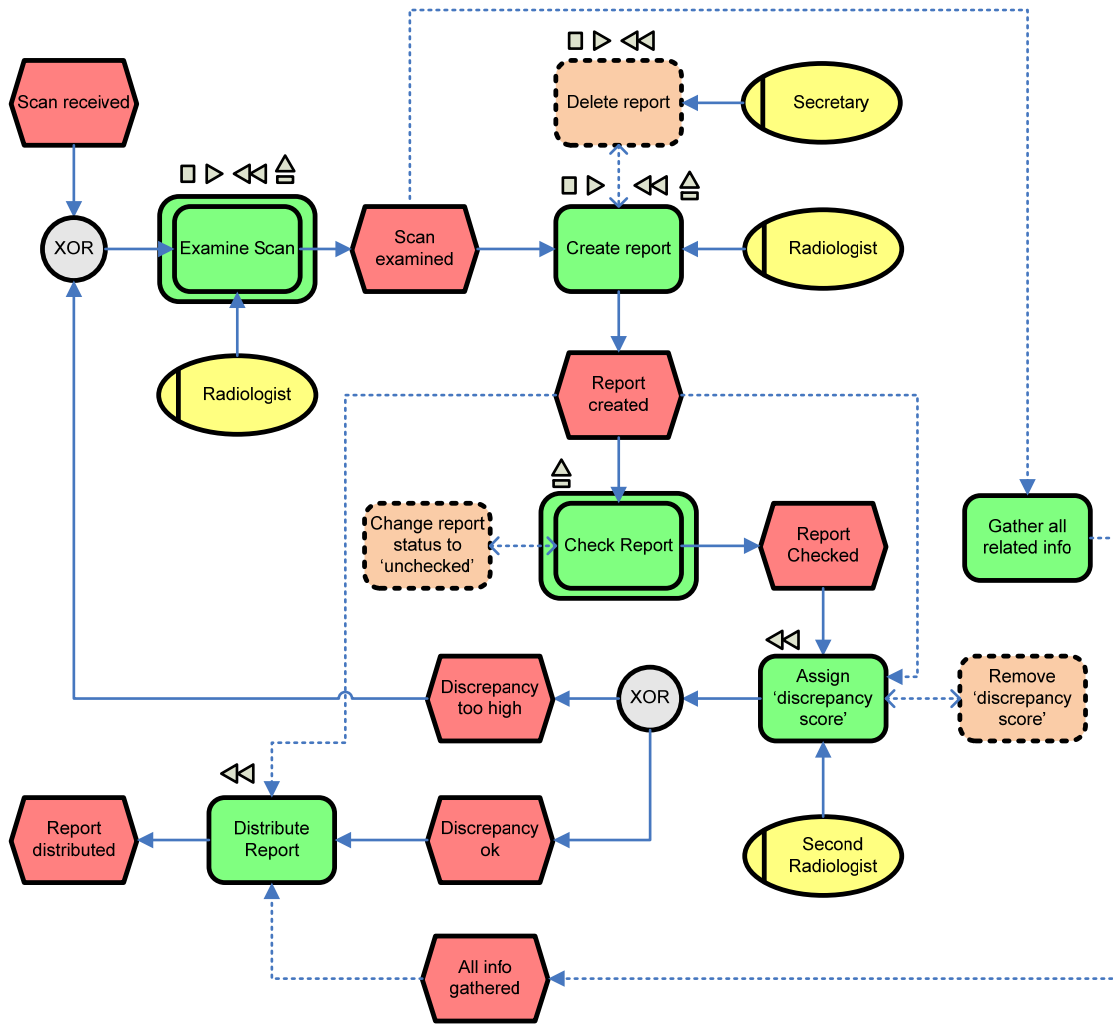


Figure 18: External view of the scan interpretation process

Appendix A.2 The Main Teleradiology Process

For a clinic or hospital that acquires medical scans (CT, MRI, X-Ray, etc) it is imperative that these scans are interpreted effectively and efficiently. If the capabilities and/or capacity is not available in-house, then outsourcing the scan interpretation, i.e., making use of a scan interpretation service, is a viable option (and maybe the only option in the future).

In such a setting, the clinic/hospital must start a search for a suitable scan interpretation service that can meet the requirements posed by the clinic/hospital. For such a search, a similar process model as shown in Figure 18 can be used. After a match has been established, an agreement can be made, possibly after negotiations, and an e-contract is established. From the e-contract the necessary system infrastructure is dynamically created at the time it is required. This infrastructure allows for the invocation of the service, but also for monitoring the progress and other agreed upon information stated in the e-contract. Figure 19 shows the main teleradiology process as it will be seen by the

clinic/hospital after 'integrating' the external view of the scan interpretation service into the main teleradiology process (excluding the transactional specifications for the non-outsourced activities).

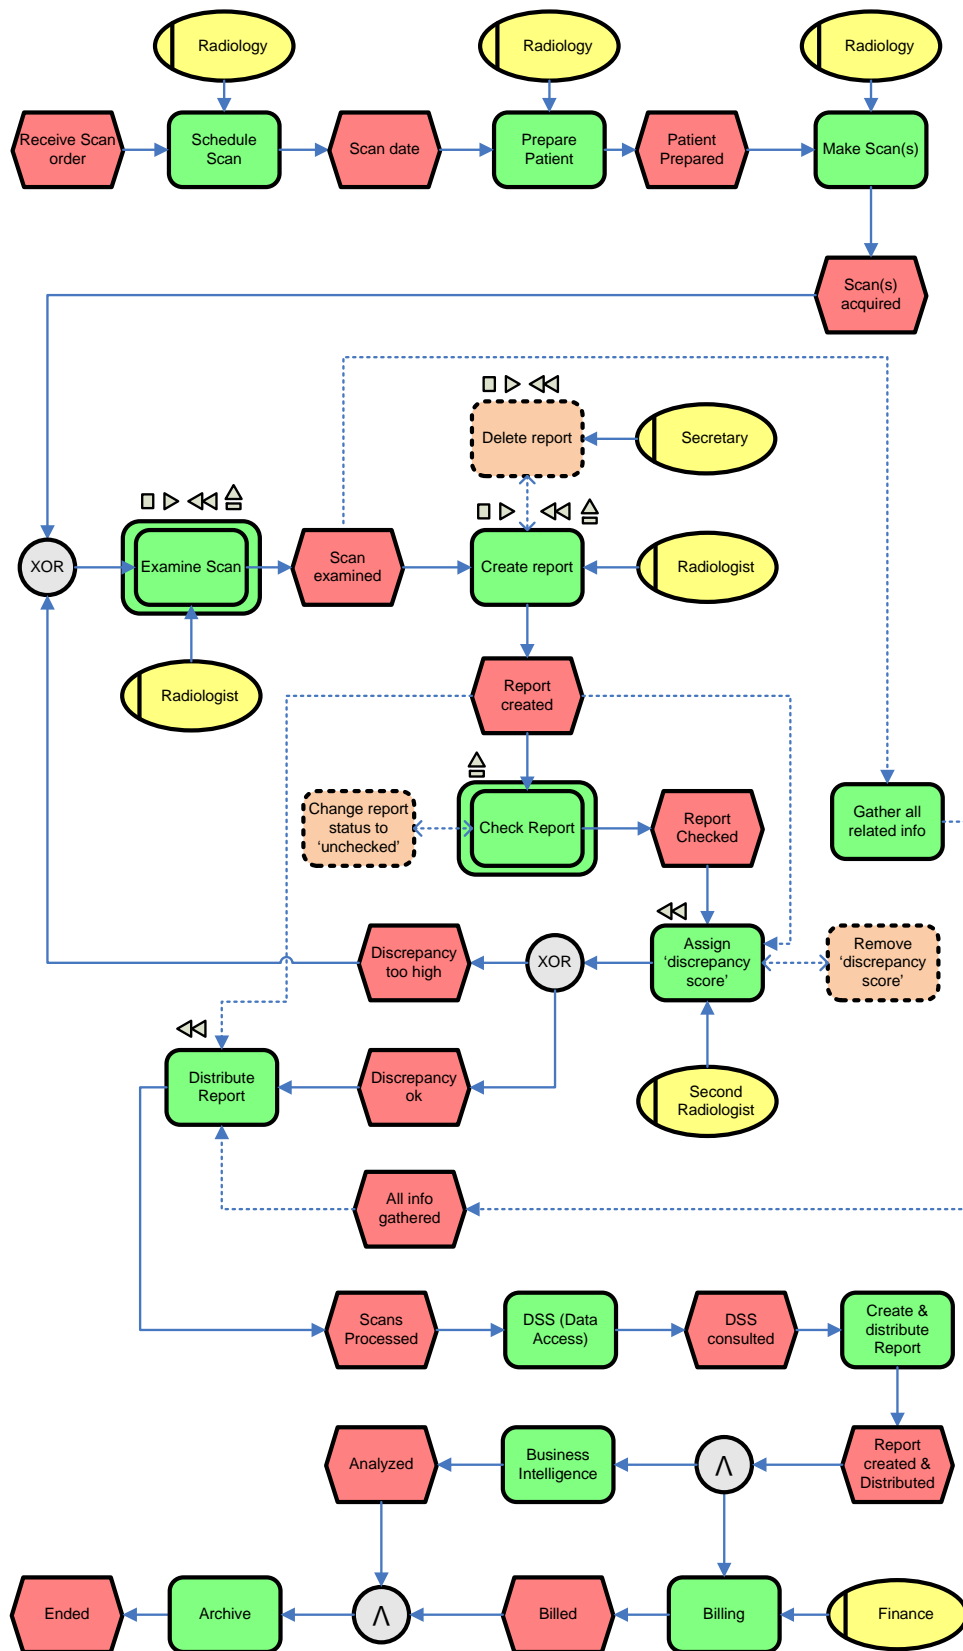


Figure 19: Global Process view